AUGUST 2013

BroadcastEngineering®

www.broadcastengineering.com

A Penton® Publication

SPECIAL REPORT:

Automation — key to efficiency and profits

ALSO INSIDE:

ABR STREAMING
A solution for poor links

FILE-BASED WORKFLOWS AND LTFS
Making LTO nearline-capable

IP INFRASTRUCTURES
Are SDI and BNCs dead?

BIG DATA
Are you ready?

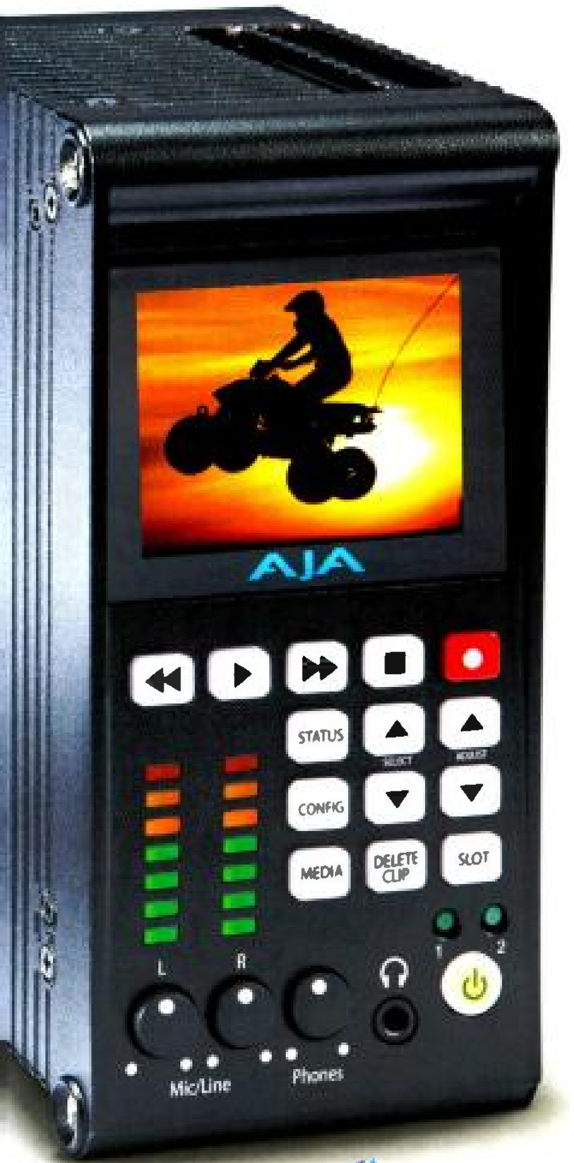

THUNDERBOLT

Ki PRO quad

Get Ready For 4K

Unite 4K, QuadHD, 2K and HD with breakthrough features and quality.

Ki Pro Quad enables breakthrough 4K, QuadHD, 2K and HD workflows combining the efficiencies of ready-to-edit Apple ProRes files and the flexibility of RAW data in a compact, solid state recording, monitoring and playback package. For the first time, a true real time lens-to-post workflow exists for any project from HD to 4K.

Ki Pro Quad records incoming HD or 4K signals as Apple ProRes files on rugged Pak SSD modules for immediate use in major editing applications. Apple ProRes files are highly efficient, saving storage space while maintaining pristine image clarity. RAW data from the Canon C500 is passed over Thunderbolt™ to compatible computer and storage systems for recording and post-processing for complete color control.

For on-set monitoring and playback, Ki Pro Quad performs a real time debayer of RAW camera data to the 4K and HD outputs, allowing you to view the native resolution and simultaneously down-converted SDI and HDMI outputs for cost-effective monitoring.

Futureproof

Work at HD resolution now with 4:4:4 quality and have the ability to move up to 4K projects using the same Ki Pro Quad.

Quality and Flexibility

Pass RAW data via Thunderbolt™ and record 4K Apple ProRes files for the perfect combination of the highest quality and manageable file sizes.

Power and Portability

Ki Pro Quad is small enough to mount directly on most cameras but powerful enough to be the hub of your 4K solution.

Find out more at www.aja.com and visit us at **Booth #7.F11**

Because it matters.™

AJA VIDEO SYSTEMS

VOLUME 55 | NUMBER 8 | AUGUST 2013

BroadcastEngineering

ON THE COVER:

Channel-in-a-box (CIB) systems have become much more reliable and include designs that allow for hybrid playout: traditional playout chains and CIBs running on one automation platform.

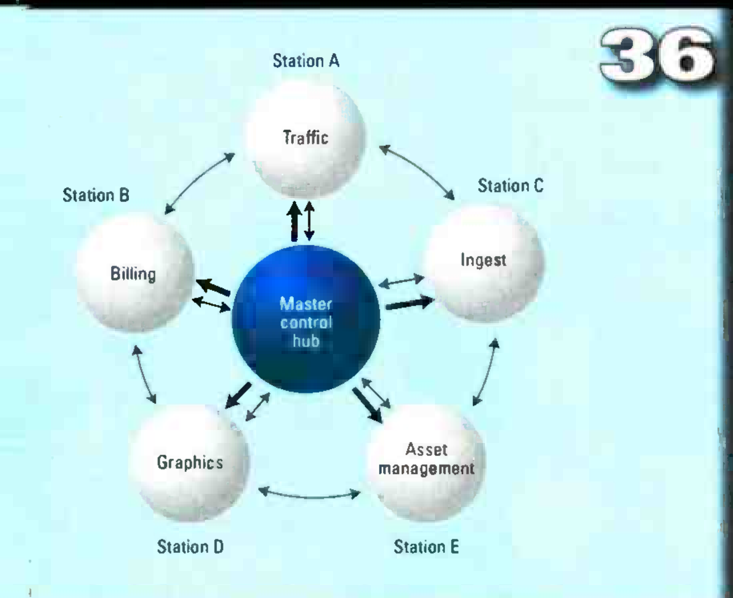

SPECIAL REPORT

BEYOND THE HEADLINES

continued on page 4

DIGITAL EDITION EXCLUSIVE FEATURES!

Check out the additional features and content found only in our digital edition:

- Special Report Feature: Metadata-driven workflows
- Field Report: Panasonic's AG-HPX600
- Field Report: Shure's Axient
- Videos
- Additional resources at the end of articles

Don't miss your chance to get this great resource delivered directly to you!

Subscribe at: onlinesuboffers.com/ginq3xca
Already a subscriber? Look for these exclusive features in your copy of our August digital edition!

TABLE OF CONTENTS (CONTINUED)

VOLUME 55 | NUMBER 8 | AUGUST 2013

14

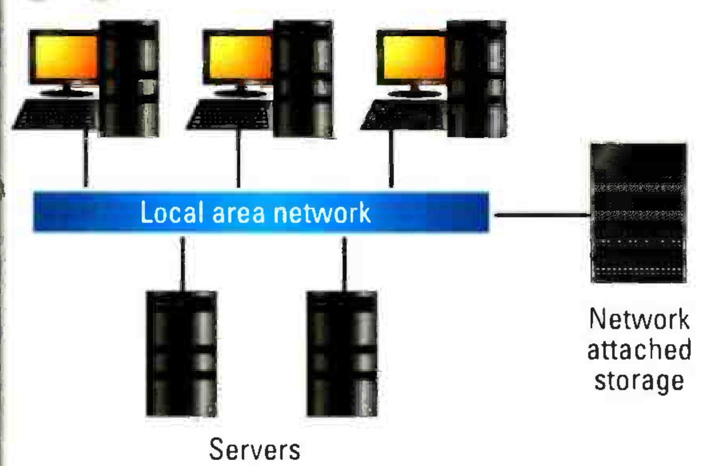

20

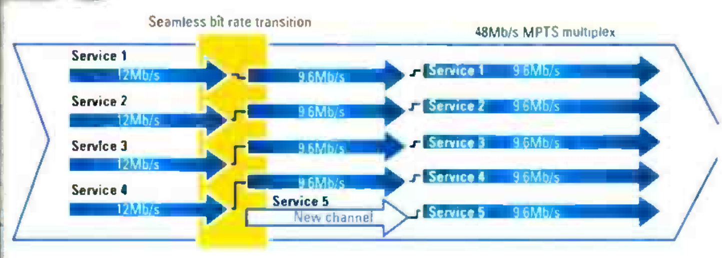

46

6

SEE IT ONLINE!

Brad On Broadcast, our editor's thoughts about the latest industry news. Read Brad's thoughts on sports, channel bundling and cable bills.

Learn more at *www.broadcastengineering.com*

If Apple can't, Intel can't

Rumors of Apple building a TV have continued for years. Sure, there's a product called "Apple TV," but it's neither a TV nor a cable box. And we have Google TV. Despite the best efforts of both companies, neither product has proven really successful in terms of delivering content. In the consumer space, being able to create and build great electronics does not guarantee a winning formula, because half of the solution is having content for those devices. Because neither Apple nor Google owns much content, and certainly not premium content, their products remain niche.

Just prior to January's CES show, *Broadcast Engineering* carried the story, "Intel set to destroy cable TV industry," which highlighted Intel's "cable-killing technology." Said Chris Davies at SlashGear, "The new push for a slice of the living room follows Intel's apparent frustration with the failure of first-gen Google TV boxes powered by its chips." According to TechCrunch, Intel has grown tired of "everyone doing a half-assed Google TV, so it's going to do it itself, and do it right."

However, within days of those reports, Intel backed off, saying the launch announcement would be delayed to Q4. The reason for the pause? It was because "obtaining licensed content proved more difficult than expected," reported *The Wall Street Journal*.

Now, some nine months later, Intel has announced it will be releasing an STB that includes live and on-demand programming. But just like Apple and Google, the chip company apparently continues to find that technology skills alone are insufficient to bring an STB + content device to market. Intel owns less content than either Apple or Google, and without content, its device is little more than a hunk of junk. Intel is learning firsthand that those who own the content, or have long-term rights to the content, are often unwilling to share the profitable pie.

In a suite near the June NCTA convention, Intel said it plans to compete with Apple, Amazon and Google with its own STB and a line-up of live and on-demand programming. To support that service, Intel was reported to be offering as much as 75 percent more that traditional cable rates for some programming.

But a recent *New York Times* article discussed the rumored practices of some existing distributors, who appear only too willing to punish wannabe players. Writing in his blog post, Richard Greenfield, an analyst at BTIG Research, said that one unnamed distributor had prevented a channel owner from selling to a service like Intel. The legality of such moves are questionable. "It most certainly is bad for consumers, as it limits competition and prevents the emergence of distributors who can provide revolutionary new ways of experiencing TV," Greenfield said.

Said the *Times* article, "Mr. Greenfield did not name any names, but several channel owners and smaller distributors said Time Warner Cable, the nation's second-largest cable company after Comcast, has been by far the most aggressive in its dealings with channels."

Responding, Glenn Britt, CEO of Time Warner said, "We may well have ones [contracts] that have that prohibition; this is not a cookie-cutter kind of business."

The Justice Department appears to have taken note of such comments and may be investigating whether the media conglomerates are violating antitrust laws by refusing to fairly negotiate with new Internet video service providers.

But do the big cable companies and content owners have anything to fear from a Justice Department that is currently neck deep in trouble with its bungling of the Fast and Furious, IRS and AP scandals? Given the department's track record, my bet is that the congloms have nothing to fear. BE

Brad Dick

EDITORIAL DIRECTOR

Send comments to: editor@broadcastengineering.com

ULTRA HD 6G SDI

Blackmagicdesign

1 Incredible Sound
Each speaker includes subwoofer for amazing sound! Hear the richness in your audio in a tiny 1 rack unit size!

2 Audio Level Meter
Big built in level meter simulates the ballistics of a standard VU audio meter in a bright easy to read design!

3 High Quality LCD
The built in display allows visibility of input video, volume level, input selection and video standard instantly!

4 Elegant Rack Design
Machined from solid aluminum, you get an elegant front panel design that looks great in your studio!

5 Advanced HDMI Monitoring
You get the latest HDMI monitoring output that supports Ultra HD! Connect to video projectors or big screen TV's!

6 Worlds First 6G-SDI
Get the latest 6G-SDI video connections for Ultra HD with a single cable! Switches between SD, HD, 2K and Ultra HD!

7 AES/EBU Digital Audio
Connect to professional audio equipment via the built in balanced AES/EBU digital audio input!

8 Analog Audio Inputs
Connect to balanced analog audio equipment and unbalanced HiFi equipment such as iPods, DVD players and more!

Get incredible quality audio monitoring with 6G-SDI and Ultra HD technology!

Blackmagic Audio Monitor gives you fantastic quality audio monitoring in a small one rack unit size! With a dual subwoofer design combined with extra wide range speakers all backed by a super powerful class A/B amplifier, Blackmagic Audio Monitor features an elegant design with big bright multi color audio meters and a built in LCD for monitoring video sources. You even get 6G-SDI input and HDMI 4K output for native 4K big screen video monitoring!

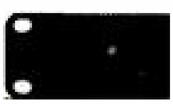

Small Size, Big Sound!
The Blackmagic Audio Monitor was designed using the latest advanced audio analysis technology to create a wide and smooth frequency response curve with crisp highs and powerful deep bass. You get great sounding music, extremely clear voice tracks and all the detail of effects tracks! Blackmagic Audio Monitor has the power to be heard in noisy environments!

Precision Audio Metering
With two big bright audio level meters on the front panel, you get full multi colored RGB LEDs under each segment so the meter scale is infinitely customizable. With lots of meter segments you can monitor a very wide dynamic range. Blackmagic Audio Monitor features an elegant and super tough front panel that's been machined from a solid block of aluminum so it looks fantastic in your studio!

Advanced 6G-SDI and Ultra HD
Blackmagic Audio Monitor supports SD, HD and the latest Ultra HD 4K video resolution via the high quality 6G-SDI input with output on the HDMI 4K connection to both HD and Ultra HD 4K displays and projectors. You can monitor audio from any of the 16 audio channels on the 6G-SDI input. 6G-SDI is four times faster than regular HD-SDI so allows image resolutions four times greater than HD! This means you get the most advanced future proof technology!

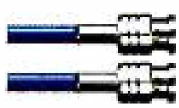

More Audio and Video Connections
Blackmagic Audio Monitor features both professional audio inputs as well as cutting edge technology video inputs including 6G-SDI input with loop output, HDMI 4K monitoring output, balanced analog XLR audio in, balanced AES/EBU audio in and consumer level HiFi audio inputs. Only Blackmagic Audio Monitor has all the inputs you need for monitoring both video and audio!

Blackmagic Audio Monitor

$1,495

Learn more today **www.blackmagicdesign.com/audiomonitor**

SONY
make.believe

Rogers Centre turns to Sony of Canada to expand and upgrade stadium control room

In operation since 1989, Rogers Centre, located in downtown Toronto, is a world-class entertainment facility and home to Major League Baseball's Toronto Blue Jays. The stadium also hosts the Canadian Football League's Toronto Argonauts as well as various events including concerts, supercross and trade shows. Rogers Centre can seat up to 67,000 spectators, depending on the event. One of the facility's standout features is its massive video board, measuring 110 feet wide and 33 feet high.

Background

Rogers Centre enhances the spectator experience through footage displayed on its main video board and smaller displays located throughout the stadium. The display footage is handled by a control room located within the facility.

For more Information on Systems Integration projects or to contact the Sony of Canada Systems Integration group please email us at psg@sony.ca

Courtesy Michelle Prata Photography/Toronto Blue Jays"

Challenge

Rogers Centre's control room had last been upgraded in 2005. All of the control room equipment was standard definition (SD) and the space was too small for the number of staff using it. By 2010, Rogers Centre staff were shooting in high definition (HD) because the content didn't just go to the stadium displays – it was also used by outside media outlets and the Toronto Blue Jays official web site. Staff were shooting in HD, then downconverting the footage so it could be played within the stadium.

The control room was also running three separate video server systems – one for replay, one for playout and one for switcher content. All three systems used different codecs and weren't networked together, so production staff were constantly having to save content to portable drives or USB keys and carry the content back and forth. The three separate systems also forced staff to create and save multiple instances of the same content for the differer codecs.

Toronto Blue Jays management wanted to upgrade its control room equipment to HD to eliminate the process of downconverting and upconverting content. They also wanted to create a larger space so video production staff could work more comfortably. Finally, they wanted a new media management and video server system that would be networked to enable staff to work more efficiently.

Case Study 034

Application: Rogers Centre
Company: Rogers Communications
Country: Canada

Solution

In 2011, Rogers Centre staff generated a rough equipment list and a basic layout and then turned to Sony of Canada who helped design and build the control room.

The space the Rogers and Sony team created is larger and more user-friendly than the old control room. The new control room supports more production staff, incorporating 24 custom-built workstations to improve operator comfort and workflow. Flexibility was improved by adding a KVM system, which allows operators to easily move their devices from one work space to another if necessary. The control room's monitor wall is now also much larger, boasting 10 42-inch Sony monitors.

The heart of the new control room is a Sony MVS-7000X switcher with 60 inputs, 48 outputs and 4 M/Es with 4 keyers per M/E. The new switcher allows production staff to work seamlessly in HD. Because the staff had already worked on other Sony switchers, they were able to learn how to use the MVS-7000X quickly.

Rogers Centre also purchased new broadcast cameras to enhance its in-house production capabilities. In addition to six existing Sony PTZ cameras, the stadium also now has five Sony HSC-300 cameras featuring high-sensitivity 2/3-inch type, 2.2 megapixel HD CCDs. These cameras are complemented by a Sony PMW-EX3 compact camcorder used to transmit footage from the stadium's stands back to the control room through a wireless backpack system. PMW-EX3 operators are often on their feet for three hours at a time, which is why a light, handheld wireless camera system was ideal.

Sony also installed a new media management and video server system from EVS Broadcast Equipment. The new system uses the ProRes LT codec and is seamlessly integrated with the control room's graphic design suites, so producers can browse all of the content in the system, make clips, create content and send that content directly back into the server system over a Gigabit Ethernet network.

Results

The new control room has now been operating smoothly for over a year with no issues. The Sony team was able to complete the control room and have everything ready for the opening of the 2012 baseball season. The switch to HD, the new workstations and the installation of the new media manage-ment system allow the workflow in the control room to run much more smoothly.

Why Sony

Rogers Centre chose to work with Sony of Canada because of the excellent reputation of Sony's Professional Solutions Integration team. The Sony team had extensive experience working with the broadcast equipment and standards Rogers Centre required.

"Working with Sony was a great experience. We knew they weren't just trying to sell us equipment. There were cases where we had suggested a certain equipment model and they found savings for us by changing to a slightly lower model. By passing on features we weren't really going to use we were able to buy more equipment. They gave us the best value possible."

– Mike Christiansen, Manager, Technical Production and Broadcast Services – Toronto Blue Jays and Rogers Centre

Products List

- MVS-7000X HD Switcher
- HSC-300 HD studio cameras
- BRCZ330 HD PTZ cameras
- PMW-EX3 camcorders
- Betacam, DVCAM, HDCAM VTRs
- PVM-1741 and PVM-2541 OLED displays
- 42-inch monitor wall displays

File-based workflows and LTFS

LTO tape can now be used as nearline storage.

BY TC CHAKRAVARTY

LTO technology was introduced to the market more than a decade ago by the LTO Consortium, comprised of the technology provider companies (TCPs), HP, IBM and Quantum. LTO is a standardized digital tape-based storage medium, and today there are hundreds of millions in use due to its ability to provide long-term shelf-stability, high capacity, reliability and high performance. Further accelerating the adoption of LTO tape is the evolution of the media and entertainment industry, which is seeking a system for managing the dramatic increases in file-based media.

However, LTO tape alone is not ideal as an archival medium for today's rapidly growing digital file-based content. Enter linear tape file system (LTFS), an open file system developed by the LTO Consortium specifically for LTO tape. It is designed to enable an open and more efficient way to access files on LTO tape. In fact, the LTO Consortium markets LTO with LTFS technology as having the ability to access files as a user would from a hard drive.

File systems and LTFS

In the past, hard-disk-based systems standardized on file systems, making it possible for applications to easily access files, and enabling hard drives to be exchangeable between systems. Examples include DOS and NTFS on Windows, HFS on Mac OS X and XFS on Linux. A disk file system defines two key components:

- The on-disk format for how content should be written to disk;
- The driver software that mounts the file system and makes it available for applications to read and write files.

These applications can be utility-based, like Finder and Windows Explorer, enabling the user to easily copy, move and rename files; or applications like Microsoft Word and Avid Media Composer, that directly operate on the files. These file systems allow any application to modify and access files, and have become a key way users access files on hard drives across one or more workstations. (See Figure 1.)

LTFS is the first open and free file system for LTO tape. While hierarchical storage management (HSM) vendors provided a file system virtualization to tape before the development of LTFS, they were not free and openly accessible. Much like disk-based systems, LTFS defines an open format for how it writes content to tape. LTFS also consists of free file system driver software that allows tapes to be mounted and be seen like any other mounted device. This allows other applications to get free and open access to LTO tape. As a result, more application providers now offer LTO with LTFS systems for media workflows.

LTFS in file-based workflows

A substantial amount of content is being archived to LTO tape in pre- and post-production environments, but LTO is typically treated as an offline, nonaccessible storage platform. In contrast, the direct access nature of LTFS enables LTO tape to be used as a nearline medium. Vast quantities of file-based media can be directly accessed from LTO tape when using LTFS workflow systems. By combining nearline and archive directly on LTO tape, this storage infrastructure solves challenges and positively impacts storage costs, data management and overall efficiencies of file-based workflows from acquisition to post. (See Figure 2.)

Here are some ways LTFS can improve file-based workflows:

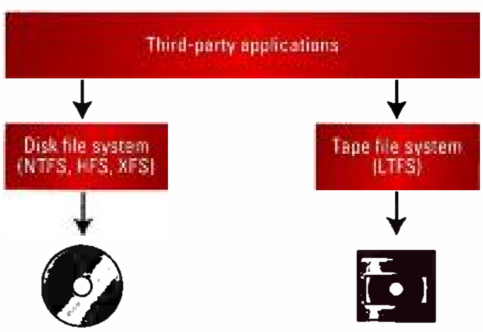

Figure 1. File systems allow third-party applications to directly access files on storage mediums.

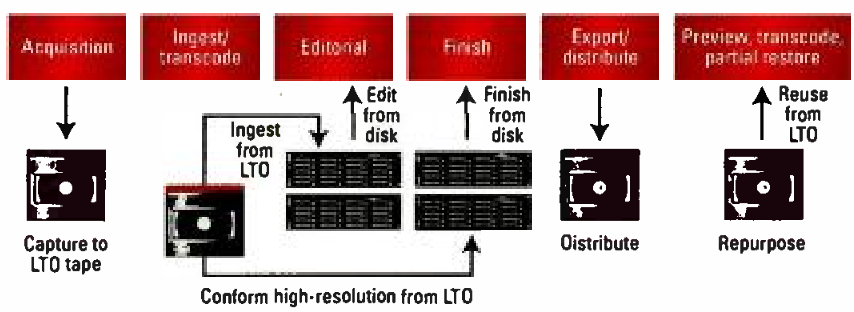

Figure 2. File-based media workflow using LTO tape with LTFS technology covers acquisition, post production, distribution and repurposing of content.


IN 5
IN 6
News from InfoComm
grass valley
Camera Integration for
MEDIORNET COMPACT

- ***Capture.*** With increasing resolution and frame rates, capturing raw file-based content has created data management challenges. Keeping large amounts of raw media on removable hard drives has its risks, and using NAS or SAN systems may be cost-prohibitive. Comparably, LTO tape is a low-cost, reliable storage medium. LTO with LTFS systems, when used in the field for quick camera master archiving, streamline workflows and create efficiencies for production crews. With LTFS, LTO has the potential to become the standard capture medium for raw, file-based media.
- ***Transport.*** Another challenge in file-based pipelines is the physical movement of large amounts of raw media. LTO with LTFS works well as both a system to capture raw media in the field and as a ruggedized transport format. Once delivered to post production, media on LTO with LTFS can also be used for post workflows.
- ***Conform.*** In videotape-based workflows, conforming was a predictable process, albeit slower. Media was redigitized from videotape at a higher resolution once the final edit was complete and was then relinked, finished and exported. Managing conforming with file-based media has become aa challenging task. Raw media is primarily located on removable drives, so the process of finding and relinking media is time consuming. However, by using LTO tape as a raw master format, it is possible to use LTFS-based systems to perform conforms directly from tape. The workflow uses an AAF, XML or EDL to identify the raw media across multiple tapes and uses LTFS to bring back just the correct files. Conforming from LTFS-formatted LTO tapes has the potential to make a challenging part of file-based workflows more automated, simpler and faster.
- ***Preview and playout.*** As broadcasters and other content owners create larger content archives on LTO tape, the ability to access this content becomes increasingly important. LTFS enables third-party applications to play files directly from LTO tape. The ability to play files for preview or playout for transmission gives media professionals a highly accessible archive that doubles as nearline storage.
- ***Transcoding and partial restore.*** LTFS-based systems allow third-party applications to directly access files on LTO tape to create dailies, low-resolution edit copies, distribution versions for Web upload or partial restores. Previously, this process was only possible from disk, but with LTFS, the raw master on LTO tape can also be used directly. Because multiple versions can be created from the archive master, media professionals can achieve quicker turnarounds on file-based workflows and better management of disk storage costs.
- ***Distribute.*** Remote WAN and file transfer systems (such as FTP and rsync) can directly access files from LTO tape for wide-area distribution. The ability to distribute directly from the archive medium without requiring restores to disk enables faster media workflows.

Conclusion

LTO tape is a proven storage medium for file-based assets but has largely been considered an offline medium, and accessibility has been a major drawback. With LTFS, this drawback has been eliminated. With LTFS, LTO tape can serve as both a nearline and archive, and LTFS-based systems enable file-based workflows directly from LTO tape. These systems have the ability to help film, video and broadcast professionals address the challenges of unprecedented content growth in the coming years. **BE**

tC Chakravarty is CEO and president, StorageDNA.

Lawo AG | Rastatt/Germany

VIDEO

by Lawo

The Lawo V__pro8 covers all of your video processing requirements in a single, compact and cost-effective 1RU package. 8x 3G/HD/SD SDI In/Out + Frame Synchronization + Variable Video & Audio Delay + 2x MADI & 1x RAVENNA In/Out + (De-) Embedding with SRC + RGB Color Correction & Proc Amp + Up/Down/Cross & AR Conversion + 5.1 Downmix + Quad-split Monitoring + Waveform & Vectorscope Display + Timecode Insertion + Test Pattern and Video ID Generator + EmBER & VSM control. And if you need to transport Video over IP in real-time, you might want to check out the new Lawo V__link4.

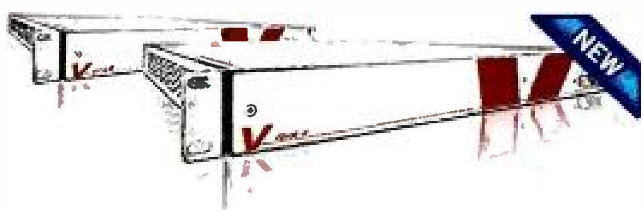

Join Lawo at IBC 2013, Hall 8, Booth 8.C71
www.lawo.ca/video

LAWO

Storage and workflow

Find your assets with proper data organization.

BY ALDO CUGNINI

Content creation and distribution have moved from a tape-only infrastructure to one that must now support multiple hard media types, including portable and fixed disks, as well as solid-state memory. All of this requires increasing amounts of mass storage, and with that, more sophisticated and complex workflows. With more content being handled in file-based form, emphasis must now be placed on file storage in the broadcast plant.

Integrated network storage

File-based data and digital content today can be stored and accessed using essentially three types of architectures, currently based on the hard-disk drive (HDD): direct attached storage (DAS), network attached storage (NAS) and storage area networks (SAN). In a DAS system, storage devices are directly connected to — and often physically integrated within — a workstation or other client system; one form of DAS is a set of hard drives within a workstation, not necessarily connected to a network.

However, when combined with a fast interface, DAS systems can provide fast data access, because each storage device specifically supports the system in which it is integrated. For this reason, DAS has long been the preferred solution in post-production.

Because DAS has the limitation that stored data cannot be shared with others for collaborative work, and it is difficult to coordinate and manage, content managers have been switching to NAS, which allows the stored content to be shared simultaneously among multiple clients. NAS storage also has the advantage of utilizing cost-effective high-speed Ethernet interconnectivity, further improving workflows.

A storage area network (SAN) is a dedicated network that connects disk storage devices directly to managing servers, using block-storage protocols; file management is left up to the controlling server. SANs have also provided a method to readily access shared content, using a specialized high-speed optical fiber or GigE networking technology. With SANs, access to data is at the block level, similar to the physical-layer storage on a hard-disk drive. SANs offer the benefit of fast, real-time content systems, but that performance comes at a financial cost. In order to provide wide access to the content, SANs were often combined with NAS, whereby the latter filled the need for near-online storage, with content archived in a lower-cost, slower-access storage tier. A SAN also requires a third-party file system to be located on each client system accessing the data on the storage network with specialized network adaptors. This makes interoperability with other, more common file systems a challenge.

In contrast, network attached storage, as shown in Figure 1, operates at the file level, providing both storage and a file system. Because of the rapid drop in the cost of building and maintaining GigE networks, as well as the increased performance level of NAS networks, many content producers and distributors are now moving toward NAS technology for content storage. In the past, in order to assure 100-percent reliability, users would rely on dual-redundant DAS or SAN servers with redundant network interfaces, so there would be no single point of failure. But one of the problems with DAS/SAN storage is scalability; when content managers need more storage, higher-speed access, or more supported users, these legacy systems most often required redundant capacity, increased complexity or completely new build-outs.

In contrast, a NAS system often can scale easily, keeping up with performance and storage needs, allowing for growth, while combining the convenience of a single expandable file system with ease of use and uncomplicated management. NAS servers ease system administration issues, improve file sharing and reliability, and provide high-speed performance. By adding support for multiple protocols, Unix and NT users can share stored files without requiring any specialized software on the clients. NAS file network protocols include FTP,

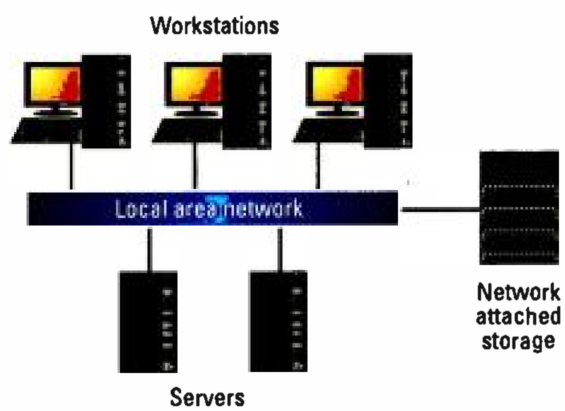

Figure 1. NAS storage consolidates workflow assets.

HTTP and NFS, providing compatibility with data management and content editing systems.

Reliability and QC

For reliability, file storage systems will often be grouped into redundant clusters, with fail-safe prioritization. One common way of doing this is in a Redundant Array of Independent Drives (RAID) storage cluster, which uses either mirroring (full redundancy) or mirroring combined with striping (distributed data) to allow all data to be accessible even if a number of drives fail. Mission-critical drive systems can also be equipped with redundant power supplies (or with battery backup) and fans, with nearly every component hot swappable, including the drives. In addition to providing drive diagnostics, health monitoring modules are available that can anticipate and head off failures; capacity expansion can be as simple as adding more drive modules to an existing system.

Solid-state drives (SSD), which have gained popularity in small laptops and netbooks, are based on rapid-access Flash memory, and would seem to be the next phase of mass-storage technology. SSD memory, which offers higher speed and reliability than the Flash memory used in USB thumb drives, could offer performance gains compared with disk-based storage.

However, SSD is probably several years from practicality in the production environment, because of high-cost, capacity limitations and unproven long-term reliability. Although a consumer-grade HDD currently runs about $100 for a 1TB (terabyte = 1000GB) 2.5in drive, an SSD of the same capacity and form factor would cost about $600. That's about $0.10/GB for the HDD and $0.60/GB for the SSD. (One hour of uncompressed 10-bit 1080i requires 585GB.)

Today's broadcast/production workflow includes many component processes. On the content/ingest side are satellite, Internet, and mobile asset sources and interfaces. On the playout side are various destinations, configured for real-time and file-based I/O that include wired and wireless interfaces. In between these sits the asset management/browsing/editing/archiving functions, which invariably rely on multiple storage solutions.

Transcoding to multiple platforms is already of growing importance, and with the ever-growing amount of content that must now be processed — much of it from amateur sources — quality control has become a critical issue. Although human input had served this function in the past, the sheer volume of production material now requires automated QC to keep up, much of it available offline, in faster-than-real-time capacity. In addition to the high-level function of detecting (and correcting when possible) file errors, sophisticated software can now detect and re-process incorrect video formats (resolutions and frame rates) and levels (skin tones, black levels). Software can now monitor audio characteristics as well, and provide a safeguard for dropouts and even the now-important loudness requirements of ATSC A/85, EBU R-128 and EBU Tech 3341.

Future storage requirements

Content handlers and producers that have not yet upgraded to state-of-the-art file storage will soon hit a critical situation that will weigh heavily on their CAPEX plans. BE

Aldo Cugnini is a consultant in the digital television industry.

ABR streaming

Learn more about this innovative way to deliver content over IP links that vary in quality over time.

BY BRAD GILMER

If you have been working with video over IP for a while, you probably are aware of at least two approaches for streaming content over IP networks. The first involves moving professional video from the venue back to the television network facility, and the other is best-effort streaming over the Internet.

The first provides highly reliable transmission over highly managed networks, probably based on the SMPTE 2022 family of standards. The second involves using video servers to stream content on the Web to end-user computers, with less than reliable results. Video freezes, audio dropouts and total loss of the feed are commonplace.

Professional transmission of video over IP is reliable, but it is relatively expensive, it requires the grooming of networks to ensure available bandwidth, and the technology is not deployed in desktop computers and smart phones. The second provides less than perfect results for people who are trying to consume broadcast content on something other than a television set. Broadcasters and media companies needed a different alternative.

ABR concept

The solution comes in the form of adaptive bit rate streaming (ABR). ABR is a technique that is used in several products, from Apple, Microsoft and others. While the specifics of the technology vary, the concept is the same. ABR technology allows the streaming of video and audio over the generic Internet without the aforementioned problems by varying the bit rate of the content being sent to the consuming device, depending upon the status of the link between sender and receiver near the time the content is to be consumed. Basically, the server sends high bit-rate content when a high bit-rate connection is available to the consumer, and it switches to lower bit-rate content when the speed of the connection is reduced.

There is a key concept behind ABR, which allows everything to fall in to place. Question: If you are receiving a stream from a server, and you initially have a great connection, but you move to a location with a poor connection, how does the server know to switch you from a higher bit-rate feed to a lower bit-rate feed? The answer is that it doesn't! It is the receiver that requests the change.

This is a key concept in ABR. ABR differs significantly from the video over IP streaming that came before it. ABR is a content pull technology, meaning that it is the receiver that requests the content, rather than the server that is pushing content to the receiver.

In a system where the receiver pulls content, the receiver knows whether it has a high bit-rate connection to the server or not. How does it know? There are several ways for the receiver to determine this. One technique is to monitor video buffer levels in the receiver. If the amount of video in the buffer is decreasing, then the bandwidth of the content is greater than the available connection. Basically, you are taking video out of the buffer and feeding it to the display faster than it is arriving at the input of the receiver. If the amount of video in the buffer is increasing, then the bandwidth of the content is lower than the available connection; the opposite condition is true. You are taking video out of the buffer and feeding it to the display slower than it is arriving at the input of the receiver.

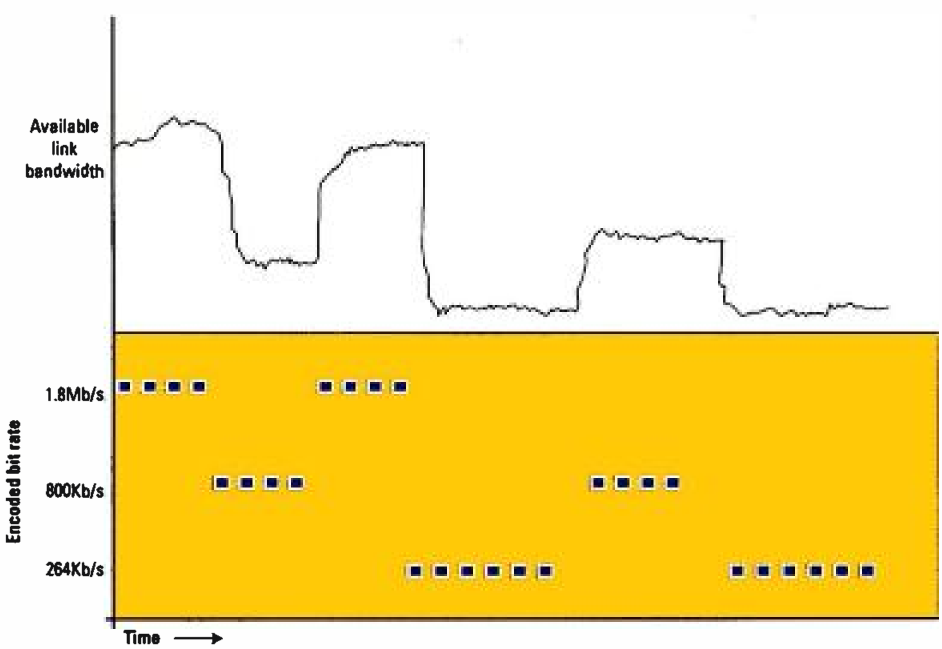

Figure 1. As available link bandwidth changes, the receiver requests content encoded at lower or higher bit rates.


AUDIO FOR TELEVISION EVOLVED
-22.5
00:07:05
DIMENSION TWO
Wheatstone

In older IP streaming systems, if the buffer in the receiver is being depleted, at some point the buffer empties and the video freezes (a buffer underflow). There is nothing the receiver can do about it.

But what if the receiver can talk with the server? What if the receiver is able to choose between several "channels," all of which contain identical content, but compressed at a different bit rate? And what if the receiver is able to ask the sender to change the "channel" being sent to it dynamically? The receiver could monitor its buffer condition. If the buffer was becoming depleted, the receiver could ask the server to send a lower bit-rate version of the same feed. If the bit rate of the content is low enough, it will arrive much more quickly than the previously streamed high bit-rate content, allowing the buffer to fill to a safe level. This approach solves the freezing problem caused by buffer underflows.

But what about moving from a low bit-rate connection to a higher one? The receiver continuously monitors the available bandwidth. Exactly how it does this varies from implementation to implementation, but in one case, the receiver probes the network by trying to download a file that is essentially bandwidth unlimited. It notes the link speed, waits for some period and then tries the download again. If the receiver notices that more bandwidth is available over a period of time, it requests a higher bit-rate "channel" from the server.

In many cases, users viewing an ABR feed are unaware that "channel" switches are being made in the background. In well-engineered ABR systems, the steps between ABR "channels" are chosen to reduce drastic changes in quality as the receiver moves from one "channel" to another. Of course, if the user is watching a feature movie on a high-quality display and the feed degrades quickly from high bandwidth to low bandwidth, then he or she is likely to see a difference. However, if the change takes place over the course of a few minutes, most users will not see the difference. In any case, a gradual loss of quality is much less noticeable to end viewers compared to a video freeze.

How exactly does this mechanism work? How are the different "channels" created? When a piece of content, a feature-length movie, for example, is ingested into an ABR system, it is typically encoded at several different bit rates. The exact bit rates depend upon the implementation, but, for our example, let's say we code at 1.8Mb/s, 800Kb/s and 264Kb/s. When the movie is ingested, three separate renditions of the movie exist on the server — one for each bit rate. This may seem like a waste of space, but it means that at any time, when the receiver requests a different bit rate, that coding rate already exists.

Chunking content

Here is another key concept of ABR systems: Content is "chunked" into many small files. A system may chop up the recorded content into three-second segments. So when the whole ingest process is finished, a single movie has been ingested and converted into a number of three-second chunks or files coded at 1.8Mb/s, the same movie has been chunked into three-second files at 800Kb/s, and again, the same movie has been chunked into three-second files at 264Kb/s.

When a receiver sees that its buffer is draining faster than it is being refilled, the receiver sends a request for the next chunk at a lower bit rate. Assuming the same network conditions, the lower bit-rate chunk can be delivered more quickly, so the receiver buffer fills more rapidly. If this lower bit-rate content can be delivered quickly enough, then the receiver continues to request three-second chunks from this lower bit-rate version until network congestion improves.

Importantly, the chunking on each file takes place at exactly the same time so that the switch between the end of one chunk of higher bit-rate content and a lower bit-rate chunk is not disruptive to the viewer. Put another way, the chunks are time-aligned across different bit-rate renditions of the movie. Chunk 1235 of the 1.8Mb/s version is *exactly* the same content as chunk 1235 of the 264Kb/s version.

Figure 1 on page 16 ties all of these concepts together. In the top portion, you can see that the available bandwidth between the server and the receiver changes over time. In the lower portion of the figure, you can see that the receiver requests higher or lower encoded bit-rate chunks depending upon available link bandwidth. This is the essence of ABR.

More to learn

There is a lot to learn about ABR technology. I hope that this introductory article whets your appetite and that you will read more about this innovative way to deliver content over IP links that vary in quality over time. BE

Brad Gilmer is executive director of the Video Services Forum, executive director of the Advanced Media Workflow Association and president of Gilmer & Associates.

? Send questions and comments to: editor@broadcastengineering.com

Seamless Channel Insertion

The technology simplifies multiplex management.

BY NICOLAS MOREAU

Television broadcasting changes quickly and is being shaped by new usages and services. Program offerings and live events access continue to grow significantly. Every day, news and sports contribution feed providers face the challenge of juggling the number of programs they must deliver and available bandwidth. As the number of programs is not constant, broadcasters and service providers need to continually optimize the use of their bandwidth on their transport multiplexes. Today, this process is primarily done by a manual intervention with limited scalability.

The MPEG standards offer an alternative way of managing bandwidth constraints dynamically named Piecewise CBR, which can be used to address this challenge and make the operational workflow as efficient as possible. This new mode allows a dynamic and seamless change of the video elementary stream bit rate. Operators in a truck or in a facility can use it to adapt and optimize the bandwidth allocated to each signal with no interruption of service. Seamless Channel Insertion (SCI) provides flexibility and opens automation possibilities in complex contribution networks.

Multifeed contribution context

Sports television program providers know well that the duration of a soccer game is not always 90 minutes, and the game may go into overtime. Tennis finals can be delayed because of rain, and a golf tournament could last longer than expected. In this situation, a new channel may need to be added in the multiplex while other programs are unexpectedly overlapping the scheduled transmission slot. Hence, the sports feed distributor could face a peak of usage in its multiplexes because the line-up of feeds to deliver becomes too important at the time. Live feed turn-around operations are often subject to immediate changes, but it is unlikely that world feed distributors will get in touch with satellite or fiber service providers to instantly adapt the bandwidth capacity. Therefore, an extra bandwidth capacity margin will be rented and saved to be used "just in case."

On the technical side, in a typical contribution multiplex, multiple-program transport stream (MPTS), each program bit rate needs to be taken into account when booking or planning the transport bandwidth capacity. This capacity will depend on the number of programs to turn around, including their specifications, such as the video format, the number of audio channels, the type of audio and the targeted video quality. Therefore, optimizing the use of the bandwidth in the multiplex can become complex.

The challenge: Avoiding service disruption when applying changes

As the program line-up could vary, the operator needs to free some bandwidth in the broadcasted multiplex to add a new channel in it. Moreover, this needs to be done without affecting the current services on-air. Contribution content distributors, such as Eurovision, from time to time, need to adjust the video bit rate of each video encoder in the same multiplexed MPTS.

Contribution video encoders configured in classic CBR need to stop and restart to produce a new MPEG-2 transport stream with a new video bit rate. This induces a stream discontinuity and therefore a "black" screen. Even if this "pause" is generally short at the encoder level, any MPEG-2 transport stream disruption will cause a disruption at the decoder level, or should we say, at the viewer level.

Considering that the new channel added in the multiplex line-up has to start at a time that may not be convenient for the other channels to be interrupted, this could lead to stressful situations for the operators.

In order to fit one or more channels in the multiplexed MPTS, the video bit rate of each current encoder needs to be reduced. The MPTS needs to remain continuous, without disruption,

Definition

Seamless Channel Insertion (SCI) is based on the open/nonproprietary standard called Piecewise CBR. Piecewise CBR extends the conventional constant-bit-rate mode (CBR) for digital video compression by allowing the change of the program bit rate on demand and without any disruption at the receiving end. Between each bit-rate change, the program is encoded at a constant bit rate. Because it relies on features offered by the MPEG open standards at both transport stream and elementary stream levels, Piecewise CBR feeds can be decoded without change using existing MPEG receiver devices.

even if a program is added or removed. Adding or removing a program in the multiplex should not impact the continuity of the other programs.

SCI rate control

Real-time video compression with MPEG-2 or H.264 technology can provide two common types of streams with regards to the bit rate control. A bit rate control algorithm in the encoder is necessary to ensure that the buffers at the encoder and decoder do not underflow or overflow[1]. Constant-bit-rate control is commonly used in the video contribution domain as it provides a constant bit rate that is mandatory for bandwidth planning and video quality assurance. Variable bit rate (VBR) is not commonly used in the contribution domain or primary distribution as it provides streams with bit rates that change, regardless of the targeted average video quality it maintains.

Still, VBR control has the interesting particularity of allowing the change of the elementary stream bit rates while maintaining the consistency of the output MPEG-2 transport stream. The SCI rate control mode is based on this VBR rate control mode, except that the video elementary stream will be changed according to the operator's request and not as a result of a video encoder algorithm.

This has no impact on the MPEG-TS multiplexing, which combines the elementary streams (video, audio channels and null packets) at the encoder level. The output is a transport stream with a constant bit rate[2], which contains a video elementary stream running at a new bit rate. Hence, the size of the MPTS at the output of the multiplexer is not changed.

Starting from four HD services running at 12Mb/s each, the operation procedure for adding a fifth channel into the multiplex using SCI would be as follows:

- The operator calculates 48Mb/s ÷ 5 = 9.6Mb/s;
- The operator assigns a new 9.6Mb/s bit rate to each encoder service by using a manual entry (typically using a dedicated Web graphic interface or a network management system that controls the encoders over SNMP). The new service bit rate is applied almost instantaneously;
- Once the four original encoders have been set to their new output bit rates, the operator can enable the fifth encoder to stream to the multiplexer. (See Figure 1.)

Starting from five HD services running at 9.6Mb/s each, a channel is removed as the program line-up changes to four services:

- The operator switches off the unwanted service;

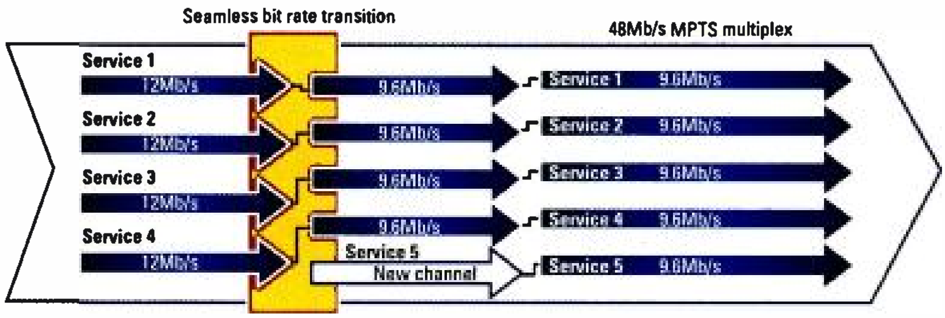

Figure 1. Example of a 48Mb/s multiplex of four services being modified to five services

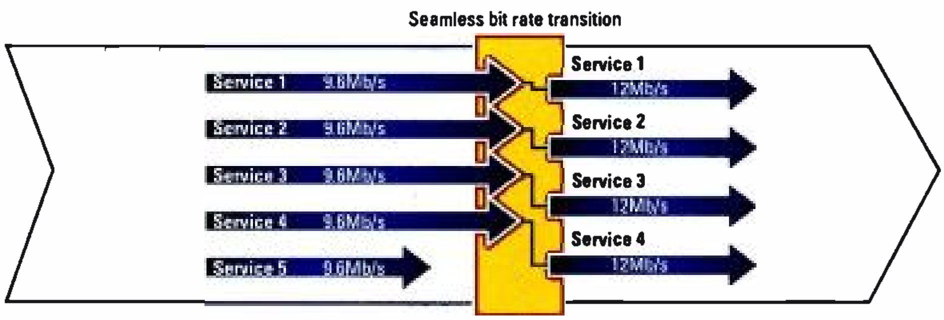

Figure 2. Example of a 48Mb/s multiplex of five services being modified to four services

Eurovision's use of Seamless Channel Insertion

Eurovision is a distributor of sports and news content for the world's broadcast and media platforms. As such, Eurovision faced the issue of its television program line-up changing regularly.

"The content broadcasted is very versatile, especially in terms of duration on-the-air," says Eurovision special project director, Puiu Dolea. "The distributed services line-up can be reshaped several times within the same day. We needed more flexibility from our encoding platform."

ATEME has implemented SCI rate control in its Kyrion line of contribution encoders at the request of Eurovision, allowing operators to change video bit rates on the fly with no service interruption. The streams delivered to EBU members and affiliates are still constant bit rate, but their bit rates can now be changed without affecting the consistency of the entire turn-around multiplex. There was no impact, no change of firmware or disruption on the receiver side following the implementation.

- The operator calculates 48Mb/s ÷ 4 = 12Mb/s;
- The operator assigns a new 12Mb/s bit rate to each encoder service. The new service bit rate is applied almost instantaneously, in this case allowing an enhanced quality for those services that remain in the multiplex. (See Figure 2.)

Control options

When SCI mode is selected, this allows the operator to set a new output bit rate at any time for each encoder service from the Web GUI or over an SNMP command. The new value will be applied almost instantaneously. Implementation into the user's interface of the encoders allows the operator to control the stream easily. Using an SNMP interface, the multiplex operator is able to integrate this new functionality into a network management system and therefore perform even more complex multiplex bandwidth management functions. By allowing operators to create a bandwidth schedule, the operator is able to know in advance the bandwidth he or she needs to book from the satellite service provider at a certain date and time.

Conclusion

Getting the most of the allocated capacity is an exciting prospect for contribution feed broadcasters. However, it could quickly become highly challenging as the layout of the television programs line-up in the MPTS can change depending on the nature of the events. SCI, an alternate rate control mode available in some MPEG-2/MPEG-4 contribution encoders, while remaining in full compliance with the standards and best practices, provides a new level of operating flexibility to broadcasters. This new technology will help contribution professionals plan and manage a feed multiplex or a turn-around feed bouquet at its best, without wasting bandwidth for the "just in case" situation that may or may not occur.

[1] Bit-Rate Control Using Piecewise, Approximated Rate-Distortion Characteristics, Liang-Jin Lin, Member, IEEE, and Antonio Ortega, Member, IEEE - IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, VOL. 8, NO. 4, AUGUST 1998
[2] ISO/IEC13818-1 - Information technology — generic coding of moving pictures and associated audio information: Systems BE

Nicolas Moreau is senior product marketing manager at ATEME and in charge of the product marketing of the Kyrion contribution encoders and decoders.

IP infrastructure

Leverage it now for more revenue later.

BY JANNE T. MORSTØL

The shift to IP within broadcasting is well under way. The multiple benefits afforded by IP networking, coupled with technology advancements and new capabilities to control quality and manage IP processes nearly to the degree you can in point-to-point networks, make this drive to IP all the more inevitable. This shift is coinciding with an astounding proliferation of connected devices that is causing a disruptive change in video and content consumption.

As we move to a converged media environment, how people consume their media will drive the market. Adopting proven IT practices and deploying IP technologies may well be the most significant path to long-term success and future revenue opportunities.

Consider IP's appeal from a core technology perspective. From a single interface, broadcasters and service and content providers can securely transport video signals to one or more receivers in a network without having to dig trenches, move cables, or add encoders or decoders at the network edge. They can also set up or add connections or take them down on demand — quickly and easily. This is powerful stuff and critical in today's broadcasting environment, which is increasingly driven by live streaming media — whether live events from a "traditional" broadcaster or streaming programming to a connected device. Broadcasters could view IP as one big virtual video router that spans the globe and runs on IP instead of traditional equipment.

Yet IP isn't just suited for the big scale. A considerable share of IP solutions in the professional video realm consist of simple networks such as studio-to-transmitter links. IP's technology allows easy and reliable connections, and when combined with Ethernet, is a universal solution for small systems to large, complex networks. IP networks are scalable, cost-effective, inherently flexible and available at a range of bandwidths, including 10Mb/s, 100Mb/s, 1Gb/s and 10Gb/s. Yet, in the broadcast world, however, IT/IP processes/signal transport must be highly protected and securely managed.

Adopting proven IT practices and deploying IP technologies may well be the most significant path to long-term success and future revenue opportunities.

Benefits by design

The advantages of IP extend beyond operational expense (OPEX) and capital expense (CAPEX) savings. Once broadcast services are managed within the IP domain, broadcasters can transform production, post-production, contribution, and distribution of core video and audio assets. The ability to share these assets quickly and efficiently on a shared IP network infrastructure can unleash unprecedented collaboration, efficiency and agility throughout the entire broadcast value chain and enable the rapid introduction of new services, including HD video and content delivery over multiple platforms.

Distribution of national and local DVB-T services to transmitter sites can also benefit from the CAPEX and OPEX saving of an IP-based network. And, once DVB-T services are managed within the IP domain, they can easily be delivered over fiber, copper or microwave networks, and within systems encompassing all three.

Finally, broadcasters can deliver TV services to consumers over multiple platforms and multiple connected screens — TV, PC and mobile devices — both in the home and on the go. Solutions must accommodate diverse video formats, quality levels and compression standards, as well as deliver the highest quality for the

Video and IP: A merger completed

Adopting proven IT practices and deploying IP technologies offers several advantages. IP technologies:

- Enable broadcasters and content providers to more easily and economically deliver content to where it is needed, when it is needed.
- Scale easily to individual needs and budgets.
- Protect investment by providing a flexible video networking system that embraces evolving architectures and services.
- Reduce operational costs and complexity through built-in intelligence.
- Ensure high-quality standards with integrated end-to-end management capabilities.
- Offer substantial OPEX savings through multiple services on a common infrastructure, with benefits multiplying as more services are migrated to IP.

True Successor of the Standard in Small HD/SD Switchers!

Enhanced Functionality with Unbelievable Cost-Performance

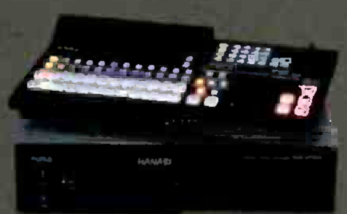

HVS-XT100OU Control Panel
HVS-XT100 Main Unit

HVS-XT110

- **HVS-XT100**, Compact 2RU frame with separate control panel
  Standard "**expandable**" I/O (8 HD/SD-SDI inputs / 4 HD/SD-SDI outputs plus 1 HDMI output)
  A. Expandable I/O (14 inputs / 6 outputs plus 1 HDMI output)
  B. Expandable I/O (12 inputs / 8 outputs plus 1 HDMI output)
  C. Variety of I/O options can be expanded in the HVS-XT100 (HD/SD-SDI, HDMI, VGA and Analog)
- **HVS-XT110**, Compact and integrated **all-in-one unit**
  Standard "fixed" I/O (12 HD/SD-SDI inputs / 6 HD/SD-SDI outputs plus 1 HDMI output)
- Both HVS-XT100/110 units come standard with **built in multi-viewer** with 4, 10, or 16 split view with tally and Title display
- **Frame Sync., Re-size engine**, and **Proc. Amp.** on inputs
- **Four keyers** are included standard;
  - Two keyers are full function with advanced chroma key
  - Two Downstream keyers (DSK)
- **Four Channels of DVE** that can be assigned to any keyer or background transitions
- **Two Channels still store** or **clip player** with 480 frames (8 seconds) of playback capability
- Future Proof **4K and 3G** "Level A" capability (TBA)

HD/SD PORTABLE VIDEO SWITCHER

HVS-XT100/XT110 NEW

HANABI

www.for-a.com

● Head Office (Japan)	Tel: +81 (0)3-3446-3936
● USA Western (CA)	Tel: +1 714-894-3311
● USA Eastern & Midwest (NJ)	Tel: +1 201-944-1120
● USA Southern (FL)	Tel. +1 305-931-1700
● Latin America & Caribbean (FL)	Tel: +1 305-931-1700
● Canada (Toronto)	Tel: +1 416-977-0343
● UK (London)	Tel: +44 (0)20-8391-7979
● Italy (Milan)	Tel: +39 02-254-3635/6
● Middle East-Africa (Dubai)	Tel: +971 4-887-6712
● Korea (Seoul)	Tel: +82 (0)2-2637-0761
● China (Beijing)	Tel: +86 (0)10-5170-9870

FOR-A for a 4K Future
Continuous Innovation

lowest cost. IP provides a common framework for easily adapting and distributing TV services for any platform or device.

"Connection-less" routing

The process of packetized content delivery, distinguished from traditional technologies such as ATM and SDH, may hold the key to IP efficiency. With traditional "connection-oriented" technologies, a path must be set up across the network from origin to destination before any traffic can be sent.

IP offers a fundamentally different paradigm, in which the network itself determines the optimal path for transmitting traffic to its destination at any given moment, and routes traffic dynamically. In the packet-based IP model, the transmission path to the destination is not set up in advance.

Instead, an end station wraps data inside a packet "container," stamps it with an origin and destination address, and sends it into the network. The network then uses the IP addresses to transport the packet to its destination through "connection-less" packet forwarding or IP routing. The nodes, or routers, forwarding these IP packets constantly update each other about the reachability of IP addresses and/or networks through the use of IP routing protocols. Today's IP routing protocols allow every router in the network to individually build a full topology view of the IP network.

This connection-less nature of IP networks provides critical advantages. Because paths are not established in advance, provisioning is easier and more cost-efficient.

IP networks are also inherently resilient; the network will always

Ethernet frame	**Ethernet, layer 2 and VLAN** Checksum, packets with error are discarded
IP header	**IP** Routing, unicast and multicast address
UDP header	**UDP** UDP port differentiates between streams; no retransmission
RTP header	**RTP** Sequence counter and timestamps
Payload MPEG-2 or FEC	**Payload** TS packets or FEC header and protection data

Table 1. The SMPTE 2022-2 standard describes how to map transport streams onto an IP network using UDP and RTP.

Connection-less IP

The connection-less nature of IP networks provides many advantages:

- The host software is much simpler at the network layer.
- The transport layer already provides connections, so no repeating is necessary.
- There is no limit to network access or traffic overload.
- An open, standards-based and widely adopted transport solution ensures future longevity and competitive pricing.
- Connection-less IP provides exceptional flexibility, with near-infinite bit-rate granularity and easily adaptable routing capabilities

reroute around any link or router failure —assuming the network has been properly designed with resilient nodes and links. This allows IP networks to survive multiple link and node failures, which is not always achieved with path-protected networking technologies. (See Table 1.)

NETWORK PERFORMANCE PARAMETER	ITU-T REC. Y.1541 QoS CLASSES	
	CLASS 6 TV CONTRIBUTION	CLASS 7 TV DISTRIBUTION
Mean IP total delay	$\leq$ 100ms	$\leq$ 400ms
IP packet delay variation	$\leq$ 50ms	
IP packet loss ratio	$\leq 1 - 10^{-5}$	
IP packet error ratio	$\leq 1 - 10^{-6}$	
Ip re-ordering ration	$\leq 1 - 10^{-6}$	

Table 2. IP network QoS class definitions and network performance objectives

More functionality, better protection

New solutions including universal IP video access modules contain all the functionality in compact form factors (as small as 1RU) that until recently required a full rack. This multi-functionality now encompasses comprehensive service monitoring, protection switching, media conversion and security, as well as management and control. This type of technology serves to increase the availability of media networks over IP, delivering quality, flexibility and end-to-end video stream protection, while solving the uncertainty of IP networks' lack of guaranteed quality, uptime, latency and packet loss.

Protection schemes including Forward Error Correction according to SMPTE 2022-1 for transport streaming, as well as diverse path routing and protection switching, and encoder protection provide the highest possible QoS and minimize the effects of random packet losses, burst packet losses, losses due to fast re-routes and link failures that can be associated with IP video transport. (See Table 2.)

Of course, end-to-end network management is required to prevent overflow by securing sufficient capacity for the media and to support controlled re-routing.

Convergence drives new opportunities

Economies of scale, built-in flexibility, lower network operating and capital costs, increased flexibility, and the ability to push more content/increased bandwidth are among IP's most compelling benefits. But perhaps even more significant is creating a reliable media network that can best serve broadcasters and content providers now and into the future.

IP is driving delivery of creative content through broadband or wireless networks for display on connected devices of all kinds, from mobile phones, laptops and notebooks to connected televisions. Business models are evolving before our eyes to profit from the growing consumer expectation for "on-demand" content.

IP networks can take advantage of new ways to monetize and leverage content — from more channels and local content insertion to the ability to push content to second screens. Greater ability to reach end users on all of their connected devices will be critical to success in the decades to come. BE

Janne T. Morstøl is chief product and marketing officer for Nevion.

? Send questions and comments to:
editor@broadcastengineering.com

ADDITIONAL RESOURCES
on the *Broadcast Engineering* website:

- IP and video merge

Special

From BroadcastEngineering

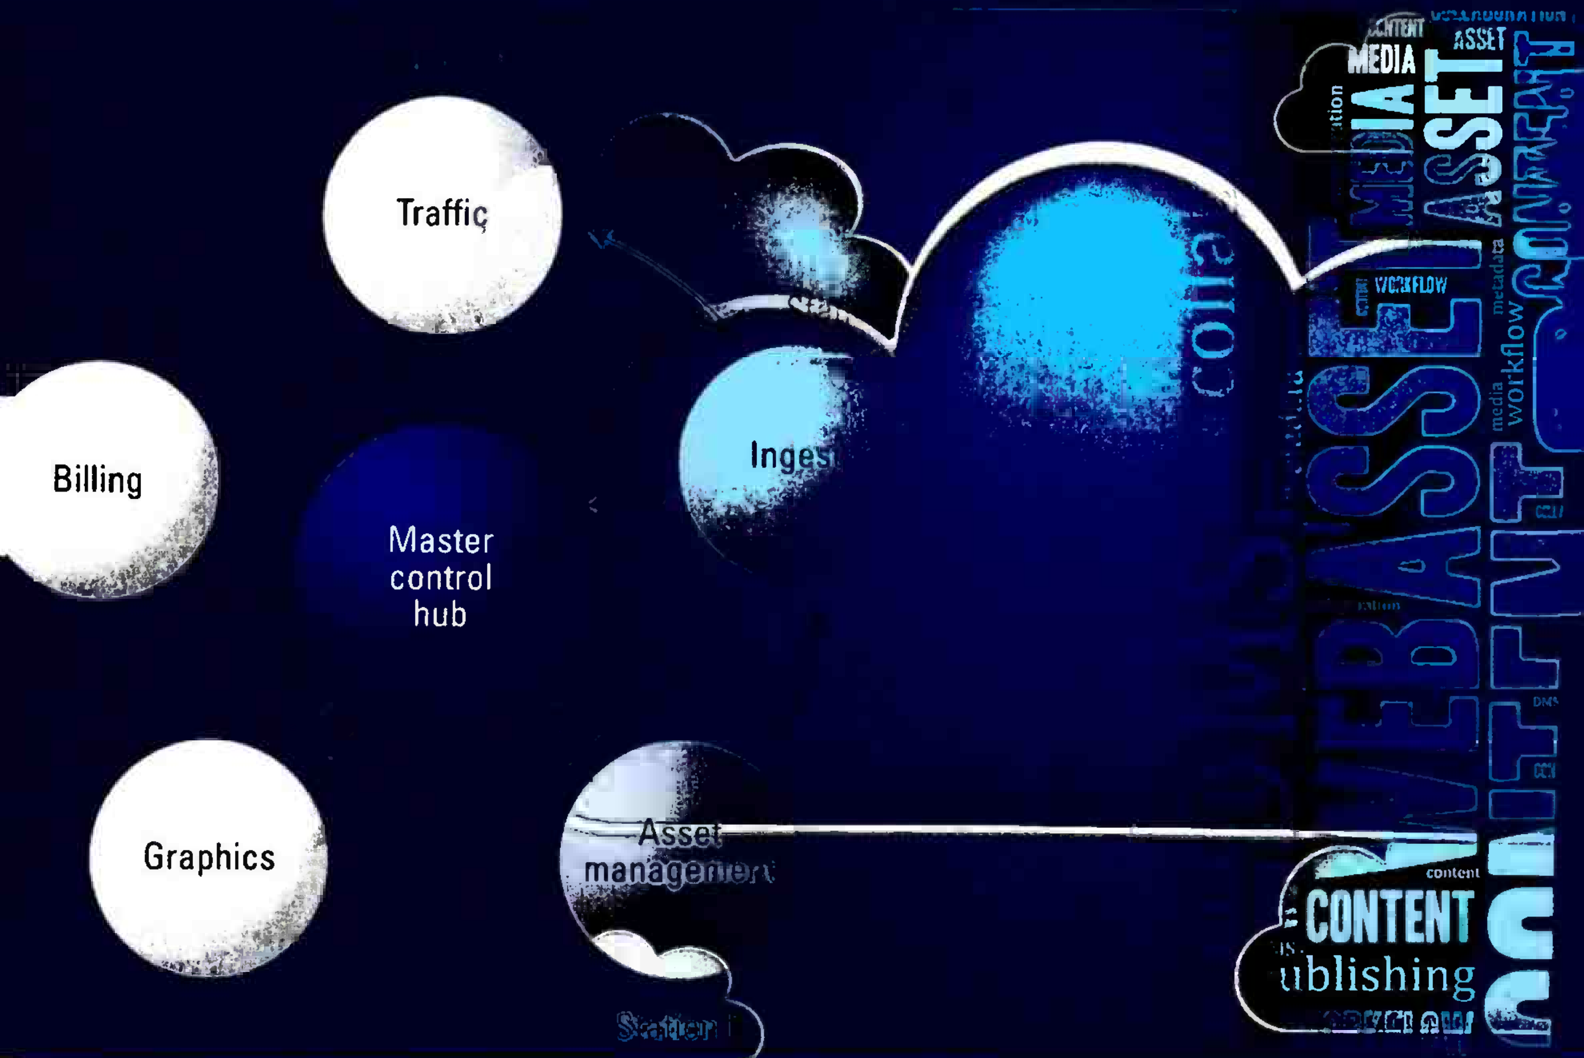

AUTOMATION

Report


metadata ASSET
WORKFLOW
LLABORATION
PUBLISHING
MEDIA
METADATA
workflow
COLLABORATION
TONIGHT ON HBTV
The Weather
up next

FLORICAL SYSTEMS®

Bridging the gap.

TV Automation & Channel-in-a-Box

Florical has been bridging the gap in TV automation since 1985. From our Channel-in-a-Box to best-of-breed or any hybrid in between - Florical is the master control solution that top broadcasters around the world have trusted for decades.

Why not have the best of both worlds by using Florical's IT centric automation with the ability to control 3rd party devices in any configuration? Florical is the only automation company that is able to provide this capability on an enterprise level. We have worked hand in hand with our clients to develop the solutions you need to elevate and secure your on air product- and of course, maximizing your revenue.

Every broadcaster is unique and your specific requirements are paramount to us. Please give us a call today, we would love to hear from you.

1.877.774.1082 FLORICAL.COM

Web, mobile, regional commercials, and narrowcasting

Are you ready for the next generation of broadcast workflows?

As "less is more" technology (commonly known as channel-in-a-box systems) come of age, broadcasters are asking, what's next? The world of traditional automation this has, in many ways, become an afterthought.

And as an automation company, it might surprise you to hear that Florical sees the next generation becoming even more transparent. "The ultimate role of automation is to become invisible" says Shawn Maynard - Florical's VP/General Manager. "The more advanced the software, the less you should know it exists." Systems should become a part of the broadcaster's workflow - not the traditional button-pushing machines. "The goal is not only to create impressive graphical user interfaces with big, beautiful buttons to click. The goal is to create a hands-free workflow that allows the content provider to do more with less, more efficiently. This is the goal of Florical - to make better solutions for our ever evolving business, and help broadcasters be more competitive."

Moving beyond the traditional roles of automation, next generation systems benefit greatly from advances in the IT/IP world. "We are moving even further towards providing our customers elegant, wholly IT- based solutions" says Dr. Eric Piard - VP of Development. "In particular, recent advancements on various technological fronts like SMPTE 2022 & 2071, the works of the VSF/EBU Task Force on Networked Media, and Software Defined Networking (SDN)." SMPTE 2022 will allow for sending real-time, uncompressed video over an IP connection - further facilitating cloud-based master control; while SMPTE 2071 enables broadcast gear to be discovered over IP - exposing its functionality and ability to be controlled. The result of the combined efforts- broadcast gear that is enterprise-wide plug and play.


Florical's AirBoss in action in Atlanta

The goal of the VSF/EBU Task Force is to standardize the production and distribution of video content - seamlessly sharing files between entities and devices. SDN separates intelligent control, normally limited to physical devices or existing software, and encourages centralizing complex tasks in both business intelligence and IT/IP gear across an entire enterprise.

One thing at the top of engineer's minds is how to effectively leverage the evolution of the cloud. The cloud's best kept secret may be what Florical has already done.

In 2007, Florical debuted the first generation of its web-based SmartCentral system. The goal was to start breaking out of traditional master control model that was task-oriented with multiple touch points, and start moving to web-based systems able to leverage private cloud for content management, content distribution, and even master control operations. "The evolution of our systems has been fun to watch" says Ash Condon- Director of Sales. "We have managed to build a system capable of acquiring content from any site, distribute that content with detailed metadata to the locations in need, all driven by demand from traffic systems, and put it on air - all without operator intervention - from any location in the world." As a TV group, why have digital acquisition equipment at every site to manage and maintain? You can acquire content once, perform quality control and file-checking once then distribute it to an infinite number of locations. The economies of scale are colossal labor savings, equipment footprints, operational costs, and overall headaches.

Reaching more viewers via additional outlets, like, web, mobile, regional commercials, and narrowcasting (targeted television advertising) are also important factors in next generation playout systems. The HD-SDI world at station and network level is moving toward IP and new codecs for 4K/8K systems and broadcasters will require more flexibility to affordably deliver content in any format across multiple platforms so they can target viewers strategically for local, regional, and national advertisers. The goal is to have access to a large amount of content across multiple screens and in multiple formats for a very low cost per channel.

The changing face of channel-in-a-box

BY TIM MENDOZA

Figure 1. Because a channel-in-a-box system has graphics, automation and DVEs, squeeze backs with "coming up" text are based 100 percent on the playlist metadata. There is no need to build and import separate graphics.

The intense discussion surrounding channel-in-a-box systems over the past 18 months has seemingly created an inaccurate impression that this is a new technology. Channel-in-a-box has actually undergone a number of permutations over the past decade, evolving from a closed architecture strategy to its current generation as a full-fledged and open, integrated channel-playout system.

Channel-in-a-Box (CIB) and integrated channel playout essentially describe the same concept of collapsing the often complex playout chain into a single server. Sometimes this server is a commodity IT server, and other times it is an OEM IT server. Some vendors rely on software for their video processing, and others rely on hardware cards.

At the start

The CIB concept was essentially born in the U.K. university town of Loughbourg, where then-CTO of OmniBus, Ian Fletcher, developed the prototype code. I constructed the prototype server in the kitchen of my Denver home using donated AMD/Microsoft components. Together, Ian and I unceremoniously unveiled the first working prototype at NAB2005. The early versions had some expected issues, notably lip sync. Those early demos included majestic footage of soaring planes and breathtaking vacation-style travel logs — anything but people talking.

Over time, this product matured. Gradually, the CIB pioneers convinced even the hardened skeptics that the future of our industry would lean heavily toward IT-based solutions. The market for CIB has grown from essentially zero in 2005 to approximately $50 million now.

The closed architecture strategy that vendors had successfully used to gain dominant market share in the early days of nonlinear editing and digital playout permitted faster development and enabled easier diagnosis of anomalies. These software issues would have been next to impossible to find in a typical complex, multi-vendor environment.

The market has since matured over several generations. CIB solutions have become much more reliable and include designs that allow for hybrid playout: traditional playout chains and CIBs running on one automation platform.

Consumer viewing habits have evolved to multiple receiving devices. Increased use of VOD services like Netflix, Hulu and Amazon signal the true arrival of the OTT market. Broadcasters simply have more options to reach the consumer as distribution channels increase and compression technologies advance. However, the resulting programming business models must work within the overall television advertising spend framework.

PricewaterhouseCoopers estimates that traditional broadcast and cable net advertising revenues will still be 90 percent of the spend in 2016. That means that 10 percent of the ad revenues must pay for the multitude of variations in new media delivery, notably mobile, Internet and OTT. This is the fertile ground

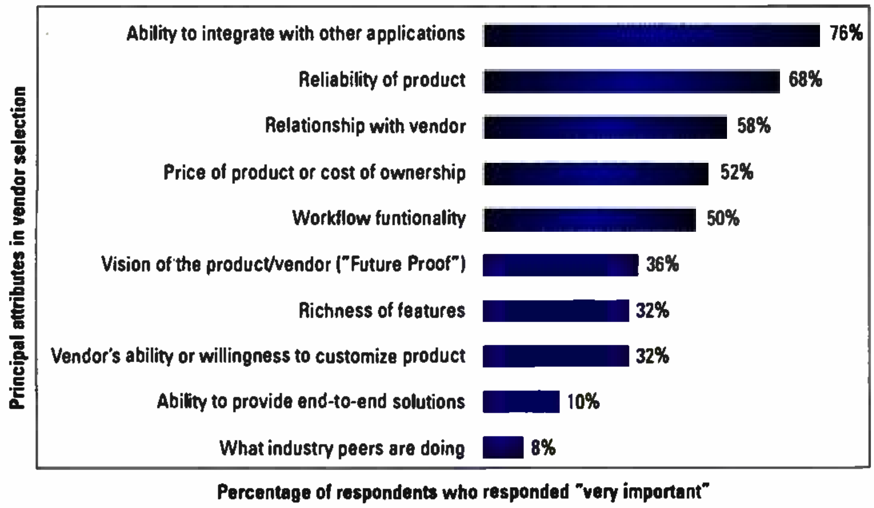

Figure 2. According to participants in a study by FrontPoint Advisors, the most important attribute in selecting a CIB vendor is the vendor's ability to integrate with other applications.

in which the CIB market will grow, as content owners are seeking ways to deliver incremental channels at minimal cost.

Today and moving forward

The CIB market encompasses not only video servers/cards and automation, but also graphics capabilities, DVEs and storage. IT servers with multiple CPU cores have taken over many of the functions that were once reserved for hardware. For example, most server products used hardware codecs for local compression for years.

Some server products took the leap with HD and went directly to software codecs. Again, many skeptical engineers were concerned due to the heavy aspect of IT functionality. This was as different codecs appeared — and, for that matter, different flavors of the same codec via simple software changes — effectively demonstrating new codec support.

Graphics initially started as simple branding and logo insertion. Having a good linear keyer was the most important aspect to good bug insertion, and today this is part of many CIB and server products.

What has evolved is taking the use of graphics to the next level — for example, having lower thirds that take dynamic text from a database or RSS feeds. L-bars are not so easy to work out, and squeeze backs traditionally take hardware-intensive digital filters to assure crisp and clear pictures, as shown in Figure 1 on page 32. The improvements in HD sets made even the tiniest artifacts show up. CIB technology stood up to this challenge with highly complex, yet flexible branding and graphics engines.

The typical CIB users have evolved over time. First adopters like satellite operators and early IPTV providers saw that the potential benefits of integrated playout far outweighed the issues associated with integrating disparate traditional hardware chains — especially when it involved hundreds or thousands of channels.

Now, CIB solutions have moved mainstream, and companies use CIB solutions not only to lower their operational expenses, but also to add the incremental channels needed to meet tight budgets. The current users of CIB solutions include not only satellite operators, but also outsource service providers, telcos and local broadcasters. Even national broadcasters are now using CIBs for second channels and, in some cases, second-screen applications.

Perhaps the greatest insight into how CIBs will evolve in the future is encapsulated in how customers view open vs. closed systems. Last year, FrontPoint Advisors conducted a broad study of the broadcast industry and found that 79 percent of respondents either preferred or were open to buying best-of-breed products and integrating the resultant solutions.

One high-profile cable network executive stated, "In the '90s and 2000s, there was significant movement toward end-to-end enterprise systems. None of the vendors could provide a solution that worked well in all areas. System integration between best-of-breed can be complex but a better solution at the end of the day." This sentiment is completely contradictory to the original CIB strategy of leveraging closed systems. Customers expect their solutions, including CIB products, to be able to integrate with other systems.

A value stack analysis from the same study also highlighted the industry's focus on discrete best-of-breed solutions. According to the survey participants, the most important attribute of selecting a vendor was the ability to integrate with other applications, as shown in Figure 2.

Business requirements of playout operations

After years of listening to such customer sentiment, I have become convinced that any such solution must integrate with and complement existing broadcaster operations. The early strategies of closed CIB solutions were right for that time, but times have evolved. In order for CIB products to grow and become an increasingly vital part of broadcasting, they must have the capability to be fully integrated with existing equipment. This means that CIB channels must co-exist in hybrid environments that contain traditional linear and new nonlinear offerings. As a system integrator succinctly stated, "Companies are rarely in a green-field situation that allows them to select an end-to-end solution from one vendor. They are usually upgrading a particular area that puts them in a market to select a vendor."

CIB should be considered as an adaptive playout solution. Breaking down playout systems into pieces that need to be adaptive, we have storage, automation, branding/graphics and, of course, file-format-independent baseband outputs. (See Figure 3.) In many cases, having a number of stand-alone devices that deliver clip playout with branding and other services is not an adequate solution. A scheme that enables blending of traditional and integrated playout channels will enable scale and simpler operations.

CIB products must have the ability to be easily integrated into the existing broadcast infrastructure. For the foreseeable future, there will be a real need for traditional, highly resilient, purpose-built traditional hardware for extremely high-value networks. In addition, there will undoubtedly be new technologies that fit a different purpose. And, of course, there will be global operations that will be tied together by a cloud-based infrastructure. The one thing that all of these models will have in common is the need to have a unified infrastructure and operational view by NOC operators.

A wise man at CBS once told me that he would never put a $3 million Super Bowl ad on a commodity IT server. I couldn't debate him, because he was right. For a $3 million Super Bowl ad, I would use the toughest traditional solution out there and back it up three times over.

That conversation has always stayed with me. There is a place for everything, and CIBs should be one of the weapons in the broadcast engineer's arsenal — not the only weapon.

Tim Mendoza is product management director at Harris Broadcast.

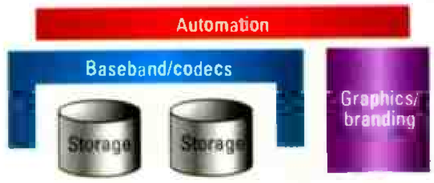

Figure 3. Adaptive playout can be one merged element or separate elements.

Centralcasting evolution

BY NED SOSEMAN

In the early 1970s, the RCA TCR-100 was 2600lbs of state-of-the-art station automation, requiring three 30A 120V circuits and engineers who knew how to fix it. Photo courtesy television.crossware.com.au

Streamlining television master control operations is nothing new. It started with the audio-follow-video master control switcher that eliminated the need for a second master control operator. Remote control panels, automated switchers and videotape cart machines promised additional savings, but in most instances those advancements actually required engineers to do more.

In the 1980s, engineers began integrating computers in master control rooms to trigger VTR and film chain remote controls. As the trend continued, custom computer programs were written to be loaded with station program logs. Some also generated as-run lists that allowed somewhat automatic client billing. Just as most technical operations systems were unique to each local TV station, nearly every station assembled its own "automation" system, often using custom TTL controls, contact closures and interfaces designed and built by local station engineers.

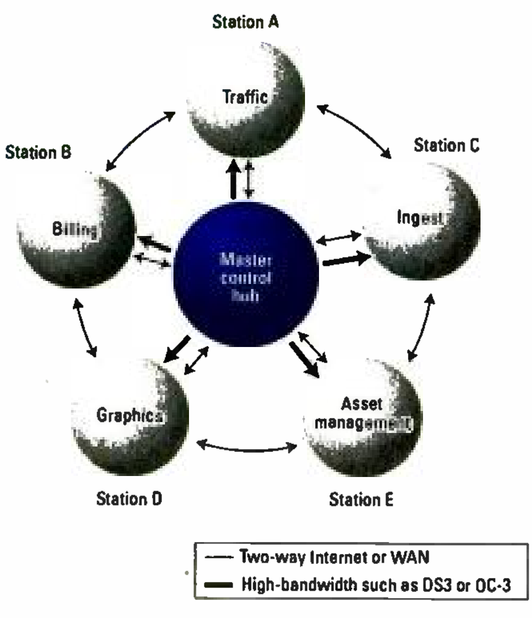

Figure 1. This hybrid centralcasting variation spreads the workflow across a group of stations.

As broadcast groups grew and WANs and the Internet became more mainstream, some forward-thinking managers began to visualize the potential for groupwide economies of scale. By the late 1990s, the now-defunct New York Times television group had begun experimenting with a new concept called centralcasting. The idea was to remotely control local equipment at group stations from a central location. The word and idea spread quickly through the industry, and many groups began investigating the pros, cons, costs and benefits of centralcasting.

Making it work

Like the many self-designed local TV station technical operations systems, broadcast groups thought up their own ideas on how to best leverage the centralcasting concept. Thus, the word centralcasting became an umbrella term covering a wide spectrum of centralized and customized cost-cutting implementations. Over time, some early implementations have been replaced by newer ideas, and some of the newest ideas are still being refined.

There's more to centralcasting than simply streaming content or centrally controlling multiple master control rooms. It provides station groups an opportunity to pool resources and reduce operating costs in a number of departments. For example, one traffic department can produce daily program logs for several stations. One business office can reconcile as-run lists and generate billing for multiple stations simultaneously. One facility can ingest and redistribute program content and spots for multiple stations. One facility can create graphics and templates. One location can provide asset management services for the group. (See Figure 1.) It's even been proven that one facility can produce and provide newscasts for multiple stations in multiple markets and time zones. The ultimate question is if these ideas can be transparent to viewers and advertisers while significantly reducing costs and boosting profits.

Models and variations

The lesson TV broadcasters learned with limited marketing agreement (LMA) experiences and multichannel video program distributors (MVPDs) is that one traffic department, one business office and one master control can successfully operate more than one channel. Centralcasting simply takes the LMA and MVPD models to the next level by reducing or eliminating group redundancy.

Successful centralcasting is a balancing act, requiring thorough identification, analysis and comparison of infrastructure requirements, capital costs, recurring costs, benefits and risks, before and after centralcast implementation. What works for one group may not work for another. Perhaps your group owns stations that are nearly all affiliated with the same national TV network and are in the same time zone. Centralized control of stations with these common denominators is far simpler than centralized control of stations affiliated with different TV networks spread across multiple time zones. This is why nearly every group has developed its own approach to centralcasting. Domestic centralcasting groups are generally adopting their own versions of typical North American models. International groups are adopting North American models and other models that are just beginning to penetrate the domestic market, such as channel-in-a-box.

Typical models in the U.S. are variations of essentially three themes. At the low end is the original New York Times model. It uses remote control of local storage and local switching, first enabled by PC-Anywhere and remotely monitored with CU-SeeMe. It's the least expensive to adapt and the most difficult to manage frame-accurately. At the high end is the hub-spoke system, sometimes referred to as the Ackerley model. It uses large servers and multichannel switchers all located at a central hub to produce complete individual program streams for regional stations and monitor by return streams, including in some instances, spot clusters for local programs such as news. Streams to spoke stations are typically delivered via OC-3 (up to 155.5Mb/s), DS3 (up to 44.7Mb/s), satellite or a combination for redundant backup. The recurring costs of high-bandwidth stream transport systems are expensive, but the model is virtually transparent and easiest to manage.

In the mid-range is the store-and-forward model, where local stations store a predetermined number of hours of forwarded content in local servers. The advantage of store-and-forward is that some content may be transferred in less than real time, perhaps via the Internet, and some may be ingested locally. This lowers bandwidth and storage requirements, keeping recurring transport stream expenses significantly lower than the Ackerley model. It can be operated from a central location, but the downside is that if the system fails, stations have a finite time to recover.

One variation of the store-and-forward model is referred to as sharecasting. In this model, the hub controls the automation, but regional stations can edit their program logs and make last-minute content changes. This systems requires only a moderate local server, local ingest and a small switcher. It requires less streaming video from the hub and can be run locally in emergencies and system failures. It also provides more local flexibility for last-moment promos, weather and news cut-ins.

Hybrid trends

A number of hybrid models are emerging as implementations and technologies mature.

One hybrid model is centralized ingest, where one location ingests virtually all program content and distributes it among stations in the group, not necessarily in real time. A variation on that idea is centralized metadata, where one location verifies segment timings and other content-related metadata and transfers the metadata to local stations that add it to matching locally ingested content, such as syndicated programs. Facilities can be operated remotely, but because most content is stored locally, bandwidth requirements and costs are significantly reduced.

Some stations are operated locally during day- and prime-time, but operated and monitored from a centralized location during overnight hours. Remote monitoring requires little bandwidth compared with 24/7 HD or UHDTV broadcast content streams.

Another model that has appeared more recently is the channel-in-a-box (CIB). This technology combines a graphics system, master control switcher, logo inserter, EAS crawls and other necessary TV elements into a single device requiring much less space, power and cooling than similar components-in-racks systems. This model is most effective when building a channel from scratch or in an extensive infrastructure overhaul.

CIB is rapidly growing in the international market because it's one of the least expensive ways to add new full-quality channels while holding capital and operating expenses to a minimum. Most CIB systems can follow the previously mentioned content storage models: local, store-forward, hub-spoke or hybrid.

Predicting the future

It is becoming clear that the technology and system models are still maturing. At its most basic level, centralcasting is simply a broad computer network, requiring care, feeding and occasional software upgrades.

Many, including SMPTE, are beginning to think cloud storage may be the ultimate solution, and it makes good sense. Redundant content storage in private and public clouds would enable stations and groups to further pool resources, continue to reduce the need for local intervention and change content distribution models for the better. The cloud business model generally assumes that the cost of more bandwidth for newer standards such as UHDTV will be offset by more competition among transport service providers.

Others predict that the industry may migrate to regionalized centralcasting, where local stations in a market are operated out of a common third-party local master control. This hybrid model could significantly reduce long-distance transport and bandwidth costs or eliminate those recurring costs with regional microwave systems while leveraging pooled resources and economies of scale.

Regardless of which model may ultimately evolve into the de facto industry standard, broadcast engineers are starting to realize that future local master control operations may require little more than a networked computer with a video card.

Ned Soseman owns LAKETV LLC in Camdenton, MO, and is chief operator at KRBK-DT, Springfield/Osage Beach, MO.

Taking advantage of cloud-based processing

BY PAOLA HOBSON

The move to the cloud — from fixed dedicated infrastructure at broadcast facilities to virtualized infrastructure located in private data centers — can offer significant efficiency and cost benefits.

It's impossible to avoid mention of "cloud" when discussing changes in broadcast workflows, yet many organizations are cautious about investing in new technologies when the benefits are not entirely proven. Whether cloud-based processing is right for an organization or a specific functional area within that organization depends on the overall workflow and environment in which that organization or functional area operates.

Small organizations keen to tightly manage costs and with limited budgets for IT overheads are usually happy with a dedicated server model, where one server or workstation is used by one person at a time for one application at a time. Even then, such organizations are already likely to be interacting with cloud-based resources, via collaborative tools such as Aspera and Signiant, albeit in a fairly limited way, focused mainly on the content transfer element of the workflow.

Increasingly, broadcasters and media organizations are using server farms or clusters, as well as virtualized resources, because, depending on the management software implemented, these allow users to gain speed advantages by running multiple jobs in parallel, assuming each server is licensed for the appropriate application.

However, a growing number of broadcast organizations have gone beyond the cluster/Virtual Machine (VM) model, and are starting to use cloud technologies, which offer dynamic and flexible means to configure processing resources. Organizations have a choice of private cloud, in which all hardware is located within their facilities (which could be on different sites) and on which runs their own cloud layer and VMs, or public cloud, where the user contracts with a service provider to access resources on-demand.

Opportunity and cost

Competition for viewers (and, therefore, subscription and advertising revenue) drives a need for more graphics processing (e.g. for branding), more visual effects (to retain viewer interest) and more content processing (e.g. re-purposing for online and VOD delivery). Timescales seem to be shorter than ever, but workflow is peaky, and organizations find that it just isn't commercially viable to buy all the hardware and software needed to meet peak demand. It would be too expensive and burdens them with expensive assets, which lose value through depreciation every day. This is especially important for post-production companies, which are reluctant to turn down clients but cannot invest in resources that may often remain unused.

Fixed infrastructure

Organizations wishing to keep resources in-house can gain efficiency savings through the use of VMs within a server farm or cluster model.

A server farm is a collection of servers, configured such that the total pool of resources is made available to software designed to take advantage of the whole system (also termed "cluster"). This allows users to run multiple jobs in parallel, assuming each server is licensed for the appropriate application. (See Figure 1.)

When an application needs to carry out some processing, server farm management software identifies a free resource and directs the execution of the processing to that hardware. Servers in the farm are generally physically co-located (e.g. in a customer's equipment room or private data center) and are connected via local area network, with access to shared content storage. A server farm can run multiple jobs in parallel, or can enable a single job to run many times faster than execution on a single machine.

The limitation of the server farm is that only applications designed for server farm operation can get the benefit of the server farm, and multiple processor licenses are needed. Not all applications can run in a parallel mode; this is only suitable where results from one element to be processed can be calculated independently of another.

Organizations also have to take into account power consumption and cooling requirements when designing their data center. In addition, backup and reconfiguration need to be considered in case of server failure. Specifically, if one server in the farm fails, the server farm needs to be reconfigured before the application can continue. There is, therefore, a need to provision for additional servers for redundancy, and to ensure that the chosen server farm management system provides automatic failover and rapid reconfiguration of the farm when needed.

Virtualization

A VM is a software representation of a computing platform, which appears to the application as a conventional processing environment, but where access to physical resources such as hard disk, network, memory or CPU are managed by the virtualization layer. The VM allows multiple instances of operating systems and applications to share the same physical resources.

Using virtualization, a single server, or groups of servers, may be set up as multiple VMs. A software layer, called the *Hypervisor*, abstracts the VMs from the underlying hardware (host server), allowing the creation of multiple VMs that can use the available hardware resources, as illustrated in Figure 2. The Hypervisor allocates resources to each VM as required, thereby allowing multiple applications to run simultaneously and efficiently.

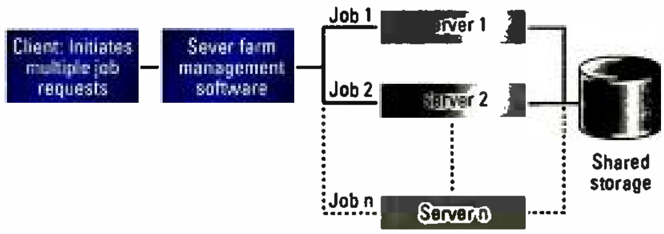

Figure 1. A server farm allows users to run multiple jobs in parallel.

VMs can offer efficiency savings because multiple VMs implemented on a server enable the hardware utilization to increase significantly. This allows users to reduce the number of servers they use, with corresponding savings in space and power.

Because VMs can also be moved, copied and reassigned between host servers, they can bring further efficiencies when a complex solution or workflow needs to be replicated. VMs also support rapid disaster recovery, where a critical application can be quickly replicated on another machine. The VMs can access centralized storage, allowing organizations to ingest the required content at the start of a project and access it from any desired application.

The main difference between the VM solution and the server farm is that the greater efficiency of the VM configuration allows broadcasters and media organizations to handle more work without necessarily having to buy more servers. In addition, because each VM appears to the application

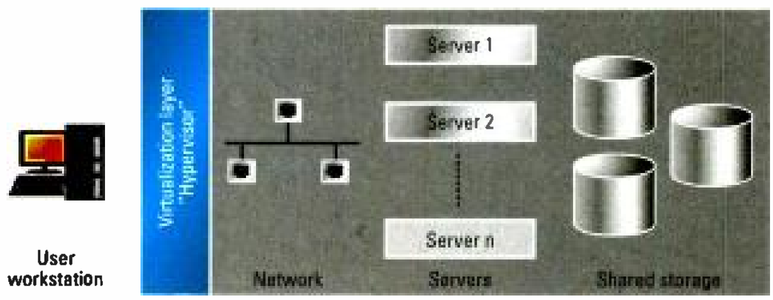

Figure 2. Using virtualization, a single server, or groups of servers, may be set up as multiple VMs. The Hypervisor abstracts the VMs from the underlying hardware (host server), allowing the creation of multiple VMs that can use the available hardware resources.

to be a conventional computer, no special applications are required, although an application license is required for each VM the application runs on.

Moving from fixed infrastructure to cloud models

For both the VM and the server farm solutions, the overall processing capability remains limited by the total amount of hardware, and neither solution solves the issue of how to manage peak demand. In addition, both require pre-configuration either into a farm configuration or into a pre-determined set of VMs. Any hardware failure requires reconfiguration and inevitable down-time.

Cloud technologies aim to solve these drawbacks by offering a dynamic and flexible means to configure processing resources. When hardware resources are augmented by a cloud layer, the system can dynamically build a VM (by assigning an available processor, memory, network connection, hard disk, etc.), collate the application image, run the required job and then release the VM resources to allow subsequent applications to run.

Organizations can set up their own private cloud, in which all hardware is located within their facilities (which could be on different sites) and on which runs their own cloud layer and VMs. Users, or automated processes, access the resources through the organization's network, as illustrated in Figure 3. Because the resources are all networked, they do not need to be physically co-located, although a high bandwidth data pipe is needed for best efficiency.

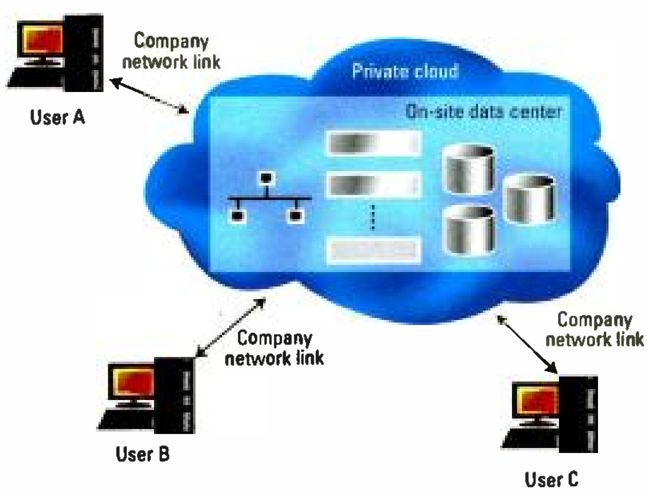

Figure 3. An advantage of using the private cloud is that all resources can be co-located.

Cloud technology vendors offer consultancy and development services to enable customers to create private cloud systems, and the organization can also outsource the system maintenance to a cloud vendor if it doesn't have the necessary skills in-house.

The advantage of the cloud model is that the technology allows resources to be assigned only when the application needs them, and only sufficient resources are allocated, which enables the entire system to be used more efficiently. The inherent flexibility allows users to configure the system to run different applications as appropriate, assuming they have available licenses. For example, a post-production company could use its system to run 10 simultaneous, high-complexity, frame rate conversions in parallel on a certain day. Then overnight, the system could reconfigure itself to support 20 simultaneous, lower complexity, format conversions.

One current drawback to the cloud solution is in finding appropriate applications that are designed to take advantage of the dynamic processing power of cloud technology. Cloud-aware applications are already available for tasks requiring massive data processing such as financial and scientific analysis. However, they are only just starting to make their appearance in the broadcast domain.

Using the private cloud

Organizations wanting to take advantage of private cloud technologies must either have skilled IT resources in-house or contract with a cloud vendor to provide managed resources. VMWare is currently a popular virtualization software in use by many broadcast organizations, although many others are in use, including Xen and KVM. Software such as VMWare provides the abstraction layer that allows a number of servers to be configured as multiple VMs, available for use when required. In a managed services model, the customer will only pay for the VMs in use at any one time. Where the private cloud is constructed within an organization's facility (private data center), there may be times when resources are not in use, which may be less efficient.

However, one advantage of the private data center is that all resources can be co-located, as shown in Figure 3, and access is via the organization's internal network, so data transfer times can be minimized. Using an open cloud technology such as OpenStack, a user can configure and manage a network of VMs. Each VM can be configured to have the operating system (e.g. Windows, Linux, etc.) required by the applications the users need to run on them, and users within the organization connect to the appropriate VM through secure log-ins. End users, e.g. operators or another application, would not need to know anything about the hardware or cloud layer; the underlying resources are hidden, and users are only aware of the application interface.

As an example, an organization wishing to take advantage of VMs for increasing transcoding efficiency would purchase multiple application licenses, so they can run the transcoder on multiple VMs. In an automated workflow, the user's media asset management system would pass the next transcoding job in the queue to an available VM, such that the processing resources in the private data center are always fully utilized.

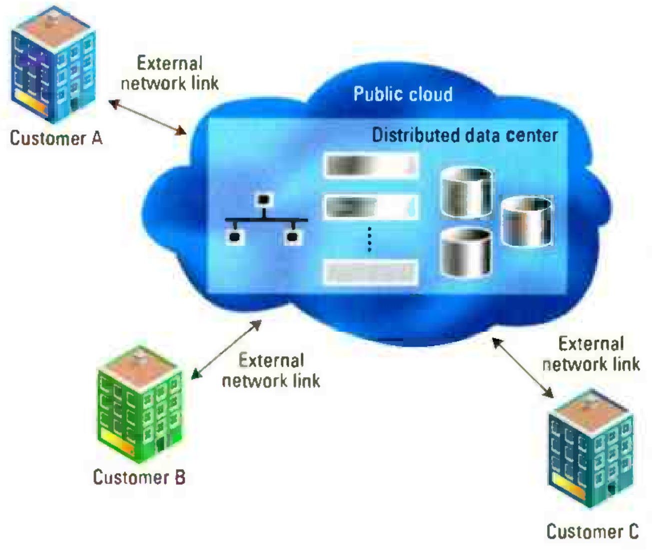

Figure 4. Using a public cloud service reduces upfront investment. Users contract with a service provider to access resources on-demand. This reduces upfront investment.

Moving capex to opex: Using the public cloud

To reduce upfront investment, organizations can take advantage of public cloud resources, which are now available from many vendors, such as Rackspace, Amazon, Salesforce and many others. (See Figure 4 on page 43.) With a public cloud service, users contract with a service provider to access resources on-demand. Customers agree on a standard level of service with a cloud provider, which enables them to carry out a number of jobs of typical size and complexity on a daily basis. This gives them access to a certain amount of storage space, memory use and processing cores.

In situations of peak demand, customers request a temporary increase in their resource allocation, which can be actioned by the cloud service provider immediately. When the peak is passed, customers return to their standard resource allocation. The advantage for customers is access to on-demand processing resources to meet peak demand, without having to invest in large amounts of hardware. They also benefit from the latest hardware and software without having to keep buying new equipment. If applications are made available as Software as a Service (SaaS), there are further advantages, because customers would only pay at execution time (pay-per-use). This level of flexible licensing allows customers to have access to a wide range of applications without requiring upfront expenditure in per-seat licenses.

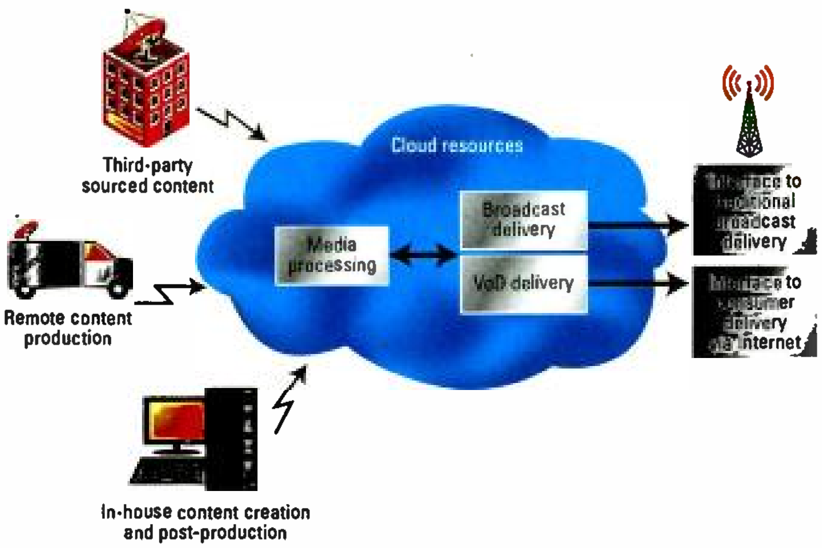

Figure 5. In this streamed media example, the cloud represents multiple servers located in different geographical regions so that the fastest and most bandwidth-efficient service can be offered to consumers worldwide.

Where the data and applications are located in a public (or private) cloud accessible internationally, organizations have the opportunity to take advantage of international subcontractors, if they wish. Because the data and applications are no longer fixed to a location, a broadcaster or media organization can instruct subsidiaries or subcontractors in any region or country to carry out the work. This can be useful for periods of peak demand, if the content owner or prime contractor does not have access to enough skilled resources locally, or if a major job comes in during a local holiday period. If the applications are such that they can be operated and managed remotely, it is possible to avoid data transfer delays when working internationally. Alternatively, if the cloud vendor has fast connections between data centers located in different countries, the required data could be transferred to wherever the workforce is available.

As an example, a broadcaster wishing to set up a streamed media service for mobile users may not initially know the scale of subscribers nor the most popular geographical locations for take-up of the service. This makes it difficult to provision resources. It may also not be clear which codec is going to be the most popular, so on-the-fly encoding is likely to be preferable to storing multiple pre-encoded copies on all the servers. In this example, a cloud model is quite advantageous. The content owner uploads the content into the cloud storage, and then the cloud vendor makes available VMs as required by consumer demand to encode and stream the material via appropriate content delivery networks. (See Figure 5.)

Note that in Figure 5, the "cloud" could actually represent multiple servers located in different geographical regions so that the fastest and most bandwidth efficient service can be offered to consumers worldwide. Cloud vendors can offer clients the use of a network of edge servers, which provide the low latency required for media services. It is also possible for the cloud vendor to offer applications such as media encoding, streaming, billing, etc. as services that are only paid for at the time of use, so the broadcaster's costs will be proportional to the number of subscribers, and the costs will only increase when revenue increases.

Cloud caveats

There are some drawbacks that must be factored into the analysis when an organization is thinking of moving from a fixed infrastructure to a cloud-based model. First, users need to consider data transfer times. Users of public cloud systems need to upload all the data required for their applications before the processing can take place, and they need to download the results. This can use up a significant amount of time, even on high bandwidth links.

Many current cloud service vendors locate their data centers in major cities, where high bandwidth connections are available, which reduces the data transfer disadvantage. Businesses also can invest in larger data pipes, thereby getting around the bottleneck, but this adds to the overall cost of the solution. Acceleration methods are also available, but these also have an associated cost.

Some organizations are reluctant to use a public cloud in case of network bottlenecks or large local simultaneous demand when they need to retrieve their data. This could arise in shared networks and in remote locations, but can be circumvented through the use of guaranteed quality of service contracts with network providers.

There also remain some fears about data security and potential data loss. However, all the major cloud storage vendors have sophisticated backup arrangements, which can, in some cases, be more secure than individual users' backup systems because multiple backups will be kept in distributed locations.

The efficiency opportunities in allowing locally, nationally or globally distributed operators to process cloud-based content assumes that suitable licensing methods have been made available by the application vendors, such that any legitimate instance of the application running anywhere within the cloud vendor's environment can access a license server. Unfortunately, this is not currently the case for many broadcast applications, where licenses are often locked to a fixed piece of hardware.

Some organizations are also worried about data security, along with rights management concerns. To address this, many media management systems, and the cloud vendors themselves, offer multiple levels of media encryption and password protection. Furthermore, some content delivery systems, such as SmartJog, include rights management layers and tracking options so content owners have all the required information.

Remember that cloud solutions currently offered by multiple vendors are not interoperable, since there are no currently agreed open standards for cloud computing. This may restrict the choices open to potential cloud users, and could make it difficult to later change cloud vendors.

Such issues will eventually disappear as standards are developed and application vendors make multiple versions to support more than one cloud technology. For example, the Cloud Data Management Interface (CDMI) being standardized under the Storage Networking Industry Association (SNIA) has defined the functional interface that applications will use to create, retrieve, update and delete data elements from the cloud. At present, this applies only to cloud storage offerings, but other standardization work is in progress within the Distributed Management Task Force (DMTF) and other groups.

Paola Hobson is senior product manager at Snell.

Data storage tools

Broadcasters can use these to store Big Data efficiently.

BY DREW ROBB

The British rock group Queen had a top-10 hit in 1989 with "I Want it All." The song included and was based around the catchphrase from singer/guitarist Brian May's second wife, "I want it all, and I want it now!"

Of course, that is what broadcast storage managers hear every day. It isn't enough to have tape archives; news editors want instant access to every inch of footage ever shot, marketing wants to know how a show is trending on Twitter as it airs, and customers demand access to music and videos that fit their personal mood at that moment.

Although Big Data analytics is one of IT's hottest topics these days, broadcast has been in the Big Data business for decades.

From a data storage aspect, this means that nothing is ever old or offline, but must be able to be ready at the push of a button. Welcome to the world of Big Data.

What is Big Data?

Big Data is a fuzzy term: It doesn't apply to any specific amount or type of data. But, as companies have created petabytes and exabytes of unstructured and semistructured data, all this data needs to be corralled, brought under control, understood and analyzed. As a rule of thumb, if there is too much data to be efficiently loaded into a relational database, it is Big Data, and specialized tools are required to turn the raw digital data into intelligence that can be used for decision making.

Although much of the emphasis is on the analytics, that is not the only aspect. For most industries, the storage and management of the data is also critical; for broadcasters, distribution also plays a huge role, given the rise of on-demand programming.

Big Data tools

Managing Big Data requires the development of a new set of tools to organize and search the data. There are now hundreds of such tools, with more on the way. Many of these are open source, or are commercial implementations of open-source software. Among the most broadly used Big Data tools are:

- ***Hadoop (hadoop.apache.org/).*** This is an open source Java-based programming framework for distributed computing of large data sets. It can scale from one server to thousands, and redistributes the work in the event of one or more nodes failing. It works with Windows, Linux, BSD and OSX.
- ***Hive (hive.apache.org).*** This is a data warehouse system that "facilitates easy data summarization, ad-hoc queries and the analysis of large datasets stored in Hadoop-compatible file systems." It is designed for batch queries/processing or large data sets and uses HiveQL, a language similar to SQL.
- ***MapReduce.*** Google developed MapReduce to address the issue of how to manage the parallel processing workloads required to process massive amounts of data. The MapReduce framework consists of two parts: Map, which distributes the load out to the compute nodes, and Reduce, which collects and merges the results from those nodes into a single result. Google researchers Jeffrey Dean and Sanjay Ghemawat published an introduction to the MapReduce framework at the Sixth Symposium on Operating Systems Development and Implementation (OSDI) in 2004, the year Google used it to replace its earlier indexing algorithms. The paper and presentation slides can be downloaded at *research.google.com/archive/mapreduce.html.*
- ***NoSQL databases.*** These are used to manage and analyze massive sets of unstructured data. NoSQL doesn't mean "no SQL," but "Not Only SQL." Structured query language (SQL) can still be used as appropriate. NoSQL databases include Apache Cassandra *(cassandra.apache.org)*, originally developed by Facebook; Amazon SimpleDB *(aws.amazon.com/simpledb/)*; and MongoDB *(www.mongodb.org/)*.

Implementing Big Data storage

Although Big Data analytics is one of IT's hottest topics these days, broadcast has been in the Big Data business for decades, having to manage media storage containing files that may be hundreds of GB or even several TB each. Big Data storage, then, takes on a different aspect than when an organization is just using Big Data for analysis. Some of the factors to consider include:

- ***Data deduplication/compression.*** A prominent feature of many storage systems is their ability to reduce the amount of disk space needed through use of deduplication or compression technologies. These technologies, however, won't have much impact on broadcasters' storage needs. Audio and video files already contain some type of compression as part of their encoding; any further compression would lose quality. In addition, there aren't as likely to be as many copies of a single video file that could be replaced with a pointer as there are with e-mails or Word documents. When there are multiple video file copies, they are likely to be needed for backup or for streaming.

- ***Tiered architecture.*** Given the breadth of data types that make up Big Data, and the different uses it is put to, a tiered storage architecture is essential. SSDs or NAND Flash memory cards can be used for analytics, storage system metadata, indexes and editing; higher-capacity, lower-cost SAS; and SATA drives for other primary and secondary storage.

Although some types of businesses are dumping tape, it can still play a key role for broadcast for setting up active archives. Video files are massive, and it can be prohibitively expensive to keep the entire archive on disks. A May 2013 TCO analysis from The Clipper Group, Inc. (*www.clipper.com/research/TCG2013009.pdf*) found that archiving on disk is 26 times as expensive as tape over a nine-year period. A tape library is slower than disk access, but each LTO-6 tape holds 2.5TB of uncompressed data, and a single Spectra Logic T950 tape library scales up to 10,200 slots, enough to hold more than 25PB of uncompressed data (62TB compressed). This allows fast, though not instantaneous, access to archives, which is good enough for archival footage.

To speed access to the videos stored in a tape library, or even lower-tiered disks, the metadata and the head of the video can be stored on the SSD drives, so access can begin immediately while the rest of the video loads from the tape.

- ***Rapid setup of multiple copies.*** Broadcast storage systems need the capacity to rapidly add and remove content from the storage being used for streaming video, scaling the number of copies to meet demand. As Joe Inzerillo, SVP and content technology/CTO for Major League Baseball's Advanced Media Group describes it, "No one cares about a 0-to-3 baseball game, until in the 7th inning, when people realize the pitcher is throwing a perfect game. Then everyone cares, and your online hits go through the roof."

- ***Common file system.*** Some editing or digital asset management systems use a proprietary storage format, which locks the company into a particular vendor and limits its ability to adopt a common storage framework for all company data. By using industry-standard methods such as NFS, CIFS and HTTP, the files can be accessed by both old and new systems, giving greater flexibility.

File system architecture has proven that it works for streaming online and on-air programming, allowing for multiple mounts for a single set of data, and it can increase throughput and capacity as needed. **BE**

Drew Robb is a freelance writer covering engineering and technology. He is author of the book "Server Management of Windows System," published by CRC Press.

ONLINE For more information on these products, go to www.broadcastengineering.com.

Grass Valley — EDIUS 7

Nonlinear video editing software is now a native 64-bit application for Windows 7 and Windows 8 64-bit operating systems, taking full advantage of up to 512GB (dependent upon Windows version) of installed memory for super-intensive media operations, such as multilayer 4K and 3D editing.

www.grassvalley.com

Leader Instruments — LF 6800

Digital broadcast signal monitor allows measurement of the channel signal corresponding to a range of globally-used broadcast signal formats; an integral Ethernet interface allows easy control of multiple LF 6800 signal monitors in any location via a single PC using a standard web browser.

www.leaderamerica.com

Apantac — MiniDE-4

CRESCENT universal mini multiviewer is a cost-effective multiviewer that displays a combination of DVI, HDMI (with HDCP), VGA, YPbPr and composite video inputs on a single display.

www.apantac.com

Clear-Com — RS-701 Beltpack

First model in a new line of RS-700 series wired analog partyline beltpacks; single-channel beltpack is equipped with an XLR-3 line connector that serves the needs of all users; combines high headroom with low-noise audio.

www.clearcom.com

RTW — BLITS

BLITS (Black & Lane's Indent Tones for Surround) instrument option has been added to the TM7 and TM9 TouchMonitors; designed to generate and analyze BLITS test signals.

www.rtw.de

LYNX Technik — yelloGUI

Mac version of yelloGUI application for yellobriks is now available; with the new yelloGUI, users have access to a host of advanced settings and adjustments of the individual yellobrik that were not possible before.

www.lynx-technik.com

Custom Consoles — Module-R Lite

Compact control desk is designed for use in production suites with limited space, or in large facilities requiring multiple rows of desks; allows easy embedding of control surfaces such as vision switches or audio mixers.

www.customconsoles.co.uk

Unified Video Technologies — uVOD

White-label solution designed to empower telcos, cable system operators (MSOs) and ISPs to quickly and easily launch state-of-the-art brand VOD services for their subscribers — without investing in costly in-house VOD technology.

www.univtec.com

Blackmagic Design — ATEM 4.1

Software update adds 1080p, 23.98, 24, 25, 29.97, 50 and 59.94 frame rates for the ATEM Production Studio 4K model; update also allows all ATEM customers to save and restore switcher states on any ATEM Production Switcher.

www.blackmagicdesign.com

CLASSIFIEDS

For Sale

Chief Engineering Job Posting:

KVOA-TV in Tucson, AZ is seeking a Chief Engineer.

Candidate must have extensive television station experience: microwave, transmitter, information systems, maintenance and building planning, combined with a strong management and planning background. Responsibilities include technical maintenance, building facilities and information systems, as well as operational and capital budget preparation and administration. This position will provide strategic planning and direction to the station's technical systems, technologies and operating standards.

Send cover letter/resume to bshaw@kvoa.com.

E.O.E.

A Penton® Publication

BroadcastEngineering

www.broadcastengineering.com

Editorial Director: Brad Dick, brad.dick@penton.com
Editor World Edition: David Austerbery, editor@broadcastengineeringworld.com
Managing Editor: Susan Anderson, susan.anderson@penton.com
Production Editor: Andrew Ward, andrew.ward@penton.com
Production Editor: Laura Collins, laura.collins@penton.com
Senior Art Director: Robin Metheny, robin.metheny@penton.com
Technical Consultants:
Computers & Networking – Brad Gilmer
Digital Video – Aldo Cugnini

Brand Manager: Carmen Lake, carmen.lake@penton.com
Sr. Marketing Manager: Kim Davidson, kim.davidson@penton.com
Director, Online Sales: Angie Gates, angie.gates@penton.com
Production Manager: Kathy Daniels, kathy.daniels@penton.com
Classified Production Manager:
Michael Penelton, michael.penelton@penton.com
Audience Development: Sonja Trent , sonja.trent@penton.com

EDITORIAL and BUSINESS OFFICE:
Penton, 9800 Metcalf Ave., Overland Park, Kansas 66212
913-341-1300, *penton.com*

MEMBER ORGANIZATIONS
Sustaining Member of: Society of Broadcast Engineers
Member, American Business Media,
Member, BPA International,
The Missouri Association of Publications

CORPORATE OFFICE:
Penton, 1166 Avenue of the Americas, 10th Floor, New York, NY 10036
Chief Executive Officer: David Kieselstein, david.kieselstein@penton.com

BE US/Canada Subscription Rates: Free/controlled circulation to qualified subscribers. Non-qualified persons may subscribe at the following yearly rates (subject to change): USA - $59; Canada - $69; Outside US and Canada - $99. Digital subscription $12 worldwide. Customer service can be reached at: *gics@pbsub.com* or by calling 866-505-7173 or 847-763-9504

BE World Subscription Rates: Free/controlled circulation to qualified subscribers. Non-qualified persons may subscribe at the following yearly rates (subject to change): USA - $59; Canada - $69; Outside US and Canada - $99. Digital subscription $12 worldwide. Customer service can be reached at: *wlcs@pbsub.com* or by calling 866-505-7173 or 847-763-9504

ARCHIVES AND MICROFORM: This magazine is available for research and retrieval of selected archived articles from leading electronic databases and online search services, including Factiva, LexisNexis and Proquest. For microform availability, contact National Archive Publishing Company at 800-521-0600 or 734-761-4700, or search the Serials in Microform listings at *napubco com.*

REPRINTS: Contact Wright's Media to purchase quality custom reprints or e-prints of articles appearing in this publication at 877-652-5295, +00-1-281-419-5725, or *penton@wrightsmedia.com*

LIST RENTAL SERVICES: MeritDirect. Contact: Marie Briganti, (877) 796-6947, mbriganti@meritdirect.com; Kristen Killingback, (914) 368-1062, kkillingback@meritdirect.com; Chris Blohm, (914) 368-1031, cblohm@meritdirect.com

PHOTOCOPIES: Authorization to photocopy articles for internal corporate, personal, or instructional use may be obtained from the Copyright Clearance Center (CCC) at 978-750-8400. Obtain further information at *copyright.com.*

PRIVACY POLICY: Your privacy is a priority to us. For a detailed policy statement about privacy and information dissemination practices related to Penton® products, please visit our Web site at *www.penton.com.*

AD INDEX

Broadcast Engineering is not responsible for errors in the Advertisers Index.

SALES OFFICES

KEY ACCOUNTS AND BRAND MANAGER

Carmen Lake
(312) 840-8473; Cell: (772) 233-2742; Fax: (913) 514-7326
carmen.lake@penton.com

US/CANADA

East
Julie Dahlstrom
(312) 840-8436
Fax: (913) 514-3684
julie.dahlstrom@penton.com

West and Classified Sales
Teri Rivas
(561) 358-6077
Fax: (772) 219-4756
teri.rivas@penton.com

Customer Service:
866-505-7173 or 847-763-9504

ISSN 0007-1994
Broadcast Engineering, August 2013, Vol. 55, No. 8 is published monthly and mailed free to qualified persons by Penton Media, Inc. 9800 Metcalf Ave., Overland Park, KS 66212-2216. Periodicals postage paid at Shawnee Mission, KS, and additional mailing offices. Canadian Post Publications Mail Agreement No. 40612608. Canada return address: IMEX Global Solutions, P.O. Box 25542, London, ON N6C 6B2. POSTMASTER: Send address changes to Broadcast Engineering, P.O. Box 2100, Skokie, IL 60076-7800 USA. CORRESPONDENCE: Editorial and Advertising: 9800 Metcalf, Overland Park, KS 66212-2216 Phone: 913-341-1300; Edit. fax: 913-967-1905. Advert. fax: 913-967-1904.

INTERNATIONAL

Europe
Richard Woolley
+44-1295-278-407
Fax: +44-1295-278-408
richardwoolley@btclick.com

Israel
Asa Talbar
Talbar Media
+972 (77) 562-1900
talbar@talbar.co.il

Japan
Mashy Yoshikawa
Orient Echo, Inc.
+81-3-3235-5961
Fax: +81-3-3235-5852
mashy@fa2.so-net.ne.jp

Asia Pacific
Tan May Ling
InterAct Media & Marketing
+65-9623-1405
Fax: +65-6265-3000
tmling@singnet.com.sg

Korea
Young J. Baek
+83-2-2273-4818
ymedia@chol.com

Beautiful video everywhere

Multiscreen is not a technology, it's an experience

Changing the way the world watches television
Amazing video quality, seamless operation, choice. Harmonic Multiscreen delivers the experience with a proven, complete solution. Compatible, tested and already in service. At scale, not in a lab. You may not see our name, but the experience will tell you that we're behind it.

Learn more at **harmonicinc.com/multiscreen**

Please visit us at **IBC**, Booth **I.B20**

harmonic